arXiv:1511.05886v2 [cs.DS] 19 Aug 2016

Randomization can be as helpful as a glimpse of the future in online computation

Jesper W. Mikkelsen*

jesperwm@imada.sdu.dk
University of Southern Denmark

Abstract

We provide simple but surprisingly useful direct product theorems for proving lower bounds on online algorithms with a limited amount of advice about the future. Intuitively, our direct product theorems say that if b bits of advice are needed to ensure a cost of at most t for some problem, then $r \cdot b$ bits of advice are needed to ensure a total cost of at most $r \cdot t$ when solving r independent instances of the problem. Using our direct product theorems, we are able to translate decades of research on randomized online algorithms to the advice complexity model. Doing so improves significantly on the previous best advice complexity lower bounds for many online problems, or provides the first known lower bounds. For example, we show that

- A paging algorithm needs $\Omega(n)$ bits of advice to achieve a competitive ratio better than $H_k = \Omega(\log k)$, where k is the cache size. Previously, it was only known that $\Omega(n)$ bits of advice were necessary to achieve a constant competitive ratio smaller than $5/4$.
- Every $O(n^{1-\varepsilon})$-competitive vertex coloring algorithm must use $\Omega(n \log n)$ bits of advice. Previously, it was only known that $\Omega(n \log n)$ bits of advice were necessary to be optimal.

For certain online problems, including the MTS, k-server, metric matching, paging, list update, and dynamic binary search tree problem, we prove that randomization and sublinear advice are equally powerful (if the underlying metric space or node set is finite). This means that several long-standing open questions regarding randomized online algorithms can be equivalently stated as questions regarding online algorithms with sublinear advice. For example, we show that there exists a deterministic $O(\log k)$-competitive k-server algorithm with sublinear advice if and only if there exists a randomized $O(\log k)$-competitive k-server algorithm without advice.

Technically, our main direct product theorem is obtained by extending an information theoretical lower bound technique due to Emek, Fraigniaud, Korman, and Rosén [ICALP'09].

*Supported in part by the Villum Foundation and the Danish Council for Independent Research, Natural Sciences.

Contents

1 Introduction **4**
- 1.1 Overview of results and techniques . . . 4
- 1.2 Lower bounds for specific online problems . . . 8
- 1.3 Relation to machine learning . . . 9
- 1.4 Other related work . . . 10
- 1.5 Further work . . . 11

2 Preliminaries **11**
- 2.1 The advice complexity model . . . 11
- 2.2 Information theoretic preliminaries . . . 13
- 2.3 Advice complexity lower bounds via Yao's principle. . . . 15

3 Direct product theorems for online computation with advice **16**
- 3.1 Definition of r-round input distributions . . . 16
- 3.2 An information-theoretic direct product theorem . . . 17
- 3.3 A martingale-theoretic direct product theorem . . . 19

4 Lower bounds for Σ-repeatable online problems **21**
- 4.1 Lower bounds for P^*_Σ . . . 22
- 4.2 Definition of repeatable online problems . . . 23
- 4.3 Proof of Theorem 1 . . . 24
- 4.4 Generalized task systems are repeatable . . . 26

5 Compact online problems **28**
- 5.1 Definition of compact online problems and proof of Theorem 2 . . . 28
- 5.2 Lazy task systems are compact . . . 29

6 Applications of main results **36**
- 6.1 Paging . . . 37
- 6.2 k-server . . . 37
- 6.3 Metrical task systems . . . 38
- 6.4 List update . . . 38
- 6.5 Bipartite matching . . . 38
- 6.6 Reordering buffer management . . . 39
- 6.7 Dynamic binary search trees . . . 39
- 6.8 Sleep states management . . . 40
- 6.9 One-dimensional unit clustering . . . 40
- 6.10 Max-SAT . . . 41
- 6.11 Metric matching . . . 41

7 Constructive version of Theorem 2 via online learning **42**

8 Repeated matrix games **43**
8.1 Proof of Theorem 3 . . . 45
8.2 Guessing strings and matching pennies . . . 46
8.3 Anti-string guessing . . . 48
8.3.1 Application to paging . . . 49
8.4 Weighted binary string guessing . . . 50
8.4.1 Proof of Theorem 4: Bin packing . . . 51

9 Superlinear advice lower bound for graph coloring **53**

10 Lower bounds for $\vee$-repeatable problems **58**
10.1 Applications . . . 60

References **61**

A Derandomization using advice **69**
A.1 Lower bounds for derandomization using advice . . . 72

B Concentration inequalities **74**

C Models of online computation **75**

D Modeling online problems as task systems **77**

E Omitted proofs **79**
E.1 Assuming same initial state . . . 79
E.2 Proof of Theorem 11 . . . 80

F Statement of main results for maximization problems **81**

1 Introduction

The model of online computation deals with optimization problems where the input arrives sequentially over time. Usually, it is assumed that an online algorithm has no knowledge of future parts of the input. While this is a natural assumption, it leaves open the possibility that a tiny amount of information about the future (which might be available in practical applications) could dramatically improve the performance guarantee of an online algorithm. Recently, the notion of advice complexity [26, 48, 63, 44] was introduced in an attempt to provide a quantitative and problem-independent framework for studying *semi-online* algorithms with limited (instead of non-existing) knowledge of the future. In this framework, the limited knowledge is modeled as a number of advice bits provided to the algorithm by an oracle (see Definition 1). The goal is to determine how much advice (measured in the length, n, of the input) is needed to achieve a certain competitive ratio. In particular, one of the most important questions is how much advice is needed to break the lower bounds for (randomized) online algorithms without advice. It has been shown that for e.g. bin packing and makespan minimization on identical machines, $O(1)$ bits of advice suffice to achieve a better competitive ratio than what is possible using only randomization [6, 3]. On the other hand, for a problem such as edge coloring, it is known that $\Omega(n)$ bits of advice are needed to achieve a competitive ratio better than that of the best deterministic online algorithm without advice [82]. However, for many online problems, determining the power of a small amount of advice has remained an open problem.

With a few notable exceptions (e.g. [24, 48, 25, 29]), most of the previous research on advice complexity has been problem specific. In this paper, we take a more complexity-theoretic approach and focus on developing techniques that are applicable to many different problems. Our main conceptual contribution is a better understanding of the connection between advice and randomization. Before explaining our results in details, we briefly review the most relevant previous work.

Standard derandomization techniques [85, 35] imply that a randomized algorithm can be converted (maintaining its competitiveness) into a deterministic algorithm with advice complexity $O(\log \log I(n) + \log n)$, where $I(n)$ is the number of inputs of length n [25] (see also Appendix A). Clearly, there are problems where even a single bit of advice is much more powerful than randomization, and so we cannot in general hope to convert an algorithm with advice into a randomized algorithm. However, using machine learning techniques, Blum and Burch have shown that a metrical task system algorithm with advice complexity $O(1)$ can be converted into a randomized algorithm without advice [20]. Problems such as paging, k-server, and list update can be modeled as metrical task systems (see e.g. [27]).

1.1 Overview of results and techniques

We will give a high-level description of the results and techniques introduced in this paper. For simplicity, we restrict ourselves to the case of minimization problems[1].

Direct product theorems. Central to our work is the concept of an *r-round input distribution*. Informally, this is a distribution over inputs that are made up of r rounds such that the requests revealed in each round are selected independently of all previous rounds. Furthermore, there must be a fixed upper bound on the length of each round (see Definition 3).

[1] All of our results (and their proofs) are easily adapted to maximization problems. See Appendix F for details.

As the main technical contribution of the paper, we prove direct product theorems for r-round input distributions. Intuitively, a direct product theorem says that if b bits of advice are needed to ensure a cost of at most t for each individual round, then rb bits of advice are needed to ensure a cost of at most rt for the entire input. This gives rise to a useful technique for proving advice complexity lower bounds. In particular, it follows that a linear number of advice bits are needed to get a (non-trivial) improvement over algorithms without any advice at all.

We provide two different direct product theorems formalizing the above intuitive statement in different ways: The first theorem (Theorem 7) is based on an information theoretical argument similar to that of [48]. The second theorem (Theorem 8) uses martingale theory.

Repeatable online problems. We combine our direct product theorems with the following very simple idea: Suppose that we have a lower bound on the competitive ratio of randomized online algorithms without advice. Often, such a lower bound is proved by constructing a hard input distribution and then appealing to Yao's principle. For some online problems, it is always possible to combine (in a meaningful way) a set of input sequences $\{\sigma_1, \ldots, \sigma_r\}$ into one long input sequence $\sigma = g(\sigma_1, \ldots, \sigma_r)$ such that serving σ essentially amounts to serving the r smaller inputs separately and adding the costs incurred for serving each of them. We say that such problems are Σ-*repeatable* (see Definition 6). For a Σ-repeatable online problem, an adversary can draw r input sequences independently at random according to some hard input distribution. This gives rise to an r-round input distribution. By our direct product theorem, an online algorithm needs linear advice (in the length of the input) to do better against this r-round input distribution than an online algorithm without advice. Thus, for Σ-repeatable online problems, we get that it is possible to translate lower bounds for randomized algorithms without advice into lower bounds for algorithms with sublinear advice. More precisely, we obtain the following theorem.

Theorem 1. *Let* P *be a* Σ*-repeatable online minimization problem and let* c *be a constant. Suppose that for every* $\varepsilon > 0$ *and every* α*, there exists an input distribution* $p_{\alpha,\varepsilon} : I \to [0,1]$ *with finite support such that* $\mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{D}(\sigma)] \geq (c-\varepsilon)\,\mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{OPT}(\sigma)] + \alpha$ *for every deterministic algorithm* D *without advice. Then, every randomized algorithm reading at most* $o(n)$ *bits of advice on inputs of length* n *has a competitive ratio of at least* c*.*

Much research has been devoted to obtaining lower bounds for randomized algorithms without advice. Theorem 1 makes it possible to translate many of these lower bounds into advice complexity lower bounds, often resulting in a significant improvement over the previous best lower bounds (see Table 1).

For a Σ-repeatable problem, the total cost has to be the sum of costs incurred in each individual round. It is also possible to consider another kind of repeatable problems, where the total cost is the *maximum cost* incurred in a single round. We call such problems $\vee$-repeatable (see Section 10). Many online coloring problems are $\vee$-repeatable. For $\vee$-repeatable problems, we show in Theorem 15 that under certain conditions, a constant lower bound on the competitive ratio of *deterministic* algorithms without advice carries over to randomized algorithms, even if the randomized algorithms have advice complexity $o(n)$. The proof of this result is straightforward and does not rely on our direct product theorems for Σ-repeatable problems. However, the result improves or simplifies a number of previously known advice complexity lower bounds for $\vee$-repeatable problems.

In Table 1, we have listed most of the repeatable online problems for which lower bounds on algorithms with sublinear advice existed prior to our work, and compared those previous lower

bounds with the lower bounds that we obtain in this paper. We have also included two examples of repeatable online problems for which Theorem 1 provides the first known advice complexity lower bounds.

It is evident from Table 1 that there are many repeatable online problems. On the other hand, let us mention that e.g. bin packing and makespan minimization are examples of problems which are provably not repeatable.

	Lower bound for algorithms with advice complexity $o(n)$			
Σ-repeatable problem	**Theorem 1**		Previous best	
Paging	$\Omega(\log k)$	[28]	5/4	[26]
k-Server	$\Omega(\log k)$	[28]	3/2	[90]
2-Server	$1 + e^{-1/2}$	[39]	3/2	[90]
List Update	3/2	[91]	15/14	[30]
Metrical Task Systems	$\Omega(\log N)$	[28]	$\Omega(\log N)$	[48]
Bipartite Matching	$e/(e-1)$	[69]	$1+\varepsilon$	[83]
Reordering Buffer Management	$\Omega(\log\log k)$	[1]	$1+\varepsilon$	[2]
2-Sleep States Management	$e/(e-1)$	[67]	7/6	[21]
Unit clustering	3/2	[49]	–	
Max-SAT	3/2	[10]	–	
$\vee$-repeatable problem	**Theorem 15**		Previous best	
Edge Coloring	2	[14]	2	[82]
$L(2,1)$-Coloring on Paths	3/2	[19]	3/2	[19]
2-Vertex-Coloring	$\omega(1)$	[16]	2	[18]

TABLE 1: New and previously best lower bounds on the competitive ratio for algorithms reading $o(n)$ bits of advice on inputs of length n. For the Σ-repeatable problems, the lower bounds in this table are obtained by combining Theorem 1 with known lower bounds for randomized online algorithms without advice. For the $\vee$-repeatable problems, the lower bounds are obtained by combining our general result on $\vee$-repeatable problems with known lower bounds for *deterministic* algorithms without advice. In both cases, references to these previously known lower bounds for online algorithms without advice are provided in the second column of the table. See Section 6 and Section 10 for a more detailed explanation of the entries in Table 1, and for a comparison with the current upper bounds.
Bipartite matching and Max-SAT are maximization problems, and hence the lower bound is obtained via Theorem 24 (which is the maximization version of Theorem 1). For paging and reordering buffer management, k denotes the cache/buffer size. For metrical task systems, N is the number of states. For the metrical task system problem and the k-server problem, the bounds are for a worst-case metric. It is also possible to use Theorem 1 together with known lower bounds for specific metric spaces. The lower bound for unit clustering is for the one-dimensional case.

Compact online problems. Note that when translating a lower bound on randomized algorithms without advice to a lower bound on algorithms with $o(n)$ bits of advice via Theorem 1, we had to make some assumptions on the lower bound. This begs the following question: Are there online problems where every lower bound (after suitable modifications) satisfies these assumptions? Recall that for finite games, Yao's principle works in both directions: In particular, if there is a lower bound for randomized computation, then there exists an input distribution witnessing this lower bound. Unfortunately, online problems are almost always infinite games (since there is no bound on the length of the input), and so this direction of Yao's principle do not necessarily apply. However, for certain online problems, it follows from previous results (see [5, 84] and Section 5.2) that this direction of Yao's principle *do* apply in the sense that the (approximately) best possible lower bound for randomized algorithms without advice is witnessed by a family of hard input distributions with finite support. We call such online problems *compact* (Definition 9). If an online problem is both Σ-repeatable and compact, we get the following result: Let c be a constant and let $\varepsilon > 0$. If c is a lower bound on the competitive ratio of randomized online algorithms without advice, then $c - \varepsilon$ is also a lower bound on the competitive ratio of online algorithms with sublinear advice. Equivalently, by contraposition, the existence of a c-competitive algorithm with sublinear advice implies the existence of a $(c + \varepsilon)$-competitive randomized algorithm without advice. Combining this with existing derandomization results [25] (see also Appendix A) yields the following complexity theoretic equivalence between randomization and sublinear advice (previously, only the forward implication was known [25]):

Theorem 2. *Let* P *be a compact and* Σ*-repeatable minimization problem with at most* $2^{n^{O(1)}}$ *inputs of length* n*, and let* c *be a constant independent of* n*. The following are equivalent:*

1. *For every* $\varepsilon > 0$*, there exists a randomized* $(c + \varepsilon)$*-competitive* P*-algorithm without advice.*
2. *For every* $\varepsilon > 0$*, there exists a deterministic* $(c + \varepsilon)$*-competitive* P*-algorithm with advice complexity* $o(n)$*.*

In this paper, we use a technique due to Ambühl [5] and Mömke [84] to show that the class of compact and Σ-repeatable problems contains all problems which can be modeled as a metrical task system (MTS) with a finite number of states and tasks. This means that e.g. the k-server, list update, paging, and dynamic binary search tree problem are all compact and Σ-repeatable, assuming that the underlying metric space or node set is finite. For all these problems, it is known that it is possible to achieve a constant competitive ratio with respect to the length of the input. Also, the number of inputs of length n for each of these problems is at most $2^{n^{O(1)}}$ (this bound holds since when we apply Theorem 2 to e.g. the k-server problem, the metric space will be fixed and not a part of the input). Thus, for each of these problems, Theorem 2 applies. Furthermore, for all the problems just mentioned (except paging), determining the best possible competitive ratio of a randomized algorithm without advice is regarded as important open problems [27, 42]. For example, the randomized k-server conjecture says that there exists a randomized $O(\log k)$-competitive k-server algorithm [74]. Theorem 2 shows that this conjecture is equivalent to the conjecture that there exists a deterministic $O(\log k)$-competitive k-server algorithm with advice complexity $o(n)$. Currently, it is only known how to achieve a competitive ratio of $O(\log k)$ using $2n$ bits of advice [25, 87].

We also show that there are compact and Σ-repeatable problems which cannot be modeled as a MTS. One such example is the metric matching problem (on finite metric spaces) [65].

Pessimistically, Theorem 2 may be seen as a barrier result which says that (for compact and Σ-repeatable online problems) designing an algorithm with sublinear advice complexity and a better competitive ratio than the currently best randomized algorithm without advice might be very difficult. Optimistically, one could hope that this equivalence might be useful in trying to narrow the gap between upper and lower bounds on the best possible competitive ratio of randomized algorithms without advice. In all cases, Theorem 2 shows that understanding better the exact power of (sublinear) advice in online computation would be very useful.

1.2 Lower bounds for specific online problems

The previously mentioned applications of our direct product theorem treat the hard input distribution as a black-box. However, it is also possible to apply our direct product theorem to explicit input distributions. Doing so yields some interesting lower bounds which cannot be obtained via Theorem 1. In what follows, we give three such examples.

Repeated matrix games. Let $q \in \mathbb{N}$ and let $A \in \mathbb{R}_{\geq 0}^{q \times q}$ be a matrix defining a two-player zero-sum game. Let V denote the value of the game defined by A. The *repeated matrix game (RMG)* with cost matrix A is an online problem where the algorithm and adversary repeatedly plays the game defined by A. The adversary is the row-player, the algorithm is the column-player, and the matrix A specifies the cost incurred by the online algorithm in each round. This generalizes the string guessing problem [24] and the generalized matching pennies problem [48] (both of these essentially corresponds to the RMG with a $q \times q$ matrix A where $A(i,j) = 1$ if $i \neq j$ and $A(i,i) = 0$). Using our direct product theorem, we easily get that for every $\varepsilon > 0$, an online algorithm which on inputs of length n is guaranteed to incur a cost of at most $(V - \varepsilon)n$ must read $\Omega(n)$ bits of advice.

Theorem 3. *Let* `ALG` *be an algorithm for the RMG with cost matrix A. Furthermore, let V be the value of the (one-shot) two-person zero-sum game defined by A and let $0 < \varepsilon \leq V$ be a constant. If $\mathbb{E}[\texttt{ALG}(\sigma)] \leq (V - \varepsilon)n$ for every input σ of length n, then* `ALG` *must read at least*

$$b \geq \frac{\varepsilon^2}{2\ln(2) \cdot \|A\|_\infty^2} n = \Omega(n) \tag{1}$$

bits of advice.

Furthermore, we also show how a more careful application of our direct product theorem to some particular repeated matrix games yields good trade-off results for the exact amount of advice needed to ensure a cost of at most αn for $0 < \alpha < V$. See Section 8 for details and for the proof of Theorem 3.

A better bin packing lower bound via repeated matrix games. We use our results on repeated matrix games to prove the following advice complexity lower bound for bin packing:

Theorem 4. *Let $c < 4 - 2\sqrt{2}$ be a constant. A randomized c-competitive bin packing algorithm must read at least $\Omega(n)$ bits of advice.*

Previously, Angelopoulos et al. showed that a bin packing algorithm with a competitive ratio of $c < 7/6$ had to use $\Omega(n)$ bits of advice by a reduction from the binary string guessing problem [6, 31]. From our results on repeated matrix games, we obtain a lower bound for *weighted* binary string

guessing. Using the same reduction as in [6], but reducing from weighted binary string guessing instead, we improve the lower bound for bin packing algorithms with sublinear advice to $4 - 2\sqrt{2}$. Thus, even though bin packing itself is not repeatable, we can obtain a better lower bound via a reduction from a repeated matrix game. The proof of Theorem 4 can be found in Section 8.4.

Superlinear lower bounds for graph coloring. We obtain the following superlinear lower bound for online graph coloring by applying our direct product theorem to an ingenious hard input distribution due to Halldórsson and Szegedy [61] (they show that a randomized graph coloring algorithm *without* advice must have a competitive ratio of at least $\Omega(n/\log^2 n)$):

Theorem 5. *Let $\varepsilon > 0$ be a constant. A randomized $O(n^{1-\varepsilon})$-competitive online graph coloring algorithm must read at least $\Omega(n \log n)$ bits of advice.*

Previously, it was only known that $\Omega(n \log n)$ bits of advice were necessary to be 1-competitive [54]. Note that $O(n \log n)$ bits of advice trivially suffice to achieve optimality for graph coloring. Furthermore, it is not hard to prove that for every $c = n^{1-o(1)}$, there exists a c-competitive graph coloring algorithm reading $o(n \log n)$ bits of advice. Thus, our lower bound saying that $\Omega(n \log n)$ bits are needed to be $O(n^{1-\varepsilon})$-competitive for every constant $\varepsilon > 0$ is essentially tight. See Section 9 for the proof of Theorem 5.

1.3 Relation to machine learning

Theorem 2 is very closely related to previous work on combining online algorithms [20, 52, 9]. Let $A_1, \ldots, A_m$ be m algorithms for a MTS of finite diameter Δ. Based on a variant of the celebrated machine learning algorithm Randomized Weighted Majority (RWM) [77], Blum and Burch obtained the following result [20, 32]: For every $\varepsilon > 0$, it is possible to combine the m algorithms into a single randomized MTS-algorithm, R, such that

$$\mathbb{E}[R(\sigma)] = (1 + 2\varepsilon) \cdot \min_{1 \leq i \leq m} A_i(\sigma) + \left(\frac{7}{6} + \frac{1}{\varepsilon}\right) \Delta \ln m, \tag{2}$$

for every input σ. An algorithm with b bits of advice corresponds to an algorithm which runs $m = 2^b$ algorithms in parallel (and selects the best one at the end). Thus, equation (2) immediately implies that given a c-competitive MTS-algorithm with advice complexity $b = O(1)$, we can convert it to a randomized $(c + \varepsilon)$-competitive algorithm without advice.

Theorem 2 improves on the result of Blum and Burch in two ways. First of all, it allows us to convert algorithms with sublinear instead of only constant advice complexity. Furthermore, Theorem 2 applies to all compact and Σ-repeatable online problems, not just those which can be modeled as a MTS.

It is natural to ask if it is possible to use the technique of Blum and Burch in order to obtain a constructive proof of Theorem 2. To this end, we remark that the result of Blum and Burch relies fundamentally on the fact that the cost incurred by RWM when switching from the state of algorithm A_i to the state of algorithm A_j for $i \neq j$ is bounded by a constant independent of n. Thus, it does not seem possible to extend the result to those compact and Σ-repeatable problems which does not satisfy this requirement (such as the metric matching problem). On the other hand, in Section 7, we show that by combining the result of Blum and Burch with the ideas that we use to prove that the MTS problem is compact, it is possible to use a variant of RWM to algorithmically

convert a c-competitive MTS-algorithm with advice complexity $o(n)$ (instead of just $O(1)$) into a randomized $(c+\varepsilon)$-competitive algorithm without advice. This yields a constructive version of Theorem 2 for problems that can be modeled as a MTS.

Very shortly after and independently of our work, Böckenhauer et al. also considered applications of machine learning techniques to advice complexity of online algorithms [22].

1.4 Other related work

Compact online problems. The existing results on (what we call) the compactness of online problems were not motivated by advice complexity. If an online problem is compact, this essentially gives a way of constructing an algorithm which can approximate the best possible competitive ratio for that online problem. Such an algorithm is known as a competitive ratio approximation scheme [57]. In general, it is not clear how construct such a scheme since online problems can have arbitrarily long input sequences (or an infinite number of possible requests in each round). For deterministic algorithms, this idea goes back at least to a paper by Manasse et al. [79]. See also [38, 78]. More recently, Ambühl showed that the list update problem is compact [5, Theorem 2.3]. Furthermore, is has been shown by Mömke in [84] that the k-server problem is compact on finite graphs. In [84], Mömke also observes that the same technique he uses to show compactness of the k-server problem can be used to show that a larger class of online problems (see [72]), including the MTS problem on finite metrics, are compact. However, [84] does not contain a full proof of this. For completeness, we show in Section 5.2 that e.g. metrical task systems with finite state spaces are compact. Our proof technique is essentially the same as the one used by both Ambühl and Mömke.

The high-entropy technique of Emek et al. As already mentioned, our information theoretical direct product theorem (Theorem 7) is an extension of a previous lower bound technique due to Emek et al. [48]. We will call this technique the *high-entropy technique*. The only known application of this technique is a lower bound for generalized matching pennies [48], but it is clear that the technique could also be used (together with an r-round input distribution) for other online problems.

Consider the set of r-round input distributions with support I. The crux of the high-entropy technique is to show that for any r-round input distribution over I, if the entropy of the distribution is sufficiently high, then every online algorithm must incur a large cost. This means that in order to achieve a low cost, an online algorithm must read enough advice so that the conditional entropy of the input sequence with respect to the advice becomes sufficiently low. In order to use the high-entropy technique, it must be the case that every probability distribution over I with high entropy yields a good lower bound. While it often happens that the uniform distribution over I (which is the distribution over I with maximum entropy) yields a good lower bound, this is far from always the case. Sometimes, good lower bounds can only be obtained by using non-uniform distributions. The main advantage of our direct product theorem over the high-entropy technique is that it can be used to obtain good lower bounds even if the uniform distribution is not a hard input distribution. In particular, this allows us to prove lower bounds for randomized online problems while treating the hard input distribution as a black-box.

1.5 Further work

We expect the techniques and results introduced in this paper to find other applications, besides those that we provide. Our main result on Σ-repeatable online problems, Theorem 1, is very easy to apply and reduces the problem of proving a linear advice complexity lower bound to that of proving a lower bound against randomized algorithms without advice. Specifically, it allows one to quickly deduce what is (implicitly) already known about the advice complexity of such problems from previous results on randomized algorithms without advice. Similarly, the results on repeated matrix games might prove useful as a starting point for showing hardness results via advice-preserving reductions.

As stated, Theorem 1 can only be used to show a constant lower bound on the competitive ratio of algorithms with sublinear advice, even if there exists a lower bound of $c(n)$ on the competitive ratio of randomized algorithms without advice, where $c(n)$ is an increasing function of the input length n. It is possible to use the techniques from this paper to obtain a more general version of Theorem 1 which given such a lower bound on randomized algorithms yields an advice complexity lower bound for achieving a competitive ratio which may depend on n. This result will appear in the final version of the paper.

2 Preliminaries

Throughout the paper, n always denotes the length of the input (number of requests). As a consequence, we use N (instead of n) to denote the number of nodes in a graph or number of points in a metric space. We usually denote an online algorithm by ALG. If we want to emphasize that an algorithm is deterministic (resp. randomized), we use the notation D (resp. R) instead. We let OPT be an optimal offline algorithm for the problem at hand. Furthermore, we let log denote the binary logarithm $\log_2$, we let ln denote the natural logarithm $\log_e$, and we let $\mathbb{R}_{\geq 0} = [0, \infty)$.

2.1 The advice complexity model

We start by formally defining competitive analysis and advice complexity. Since we are very much interested in sublinear advice, we will use the *advice-on-tape model* [26, 63]. In this model, the algorithm is allowed to read an arbitrary number of advice bits from an advice tape. There is an alternative model, the *advice-with-request model* [48], where a fixed number of advice bits is provided along with each request.

We give the formal definition of the advice-on-tape model and competitive analysis only for minimization problems, but it can easily be adapted to maximization problems.

Definition 1 (Advice complexity [26, 63] and competitive ratio [68, 88])**.** The input to an online problem, P, is a sequence $\sigma = (s, x_1, \ldots, x_n)$. We say that s is the *initial state* and $x_1, \ldots, x_n$ are the *requests*. A *deterministic online algorithm with advice*, ALG, computes the output $\gamma = (y_1, \ldots, y_n)$, under the constraint that y_i is computed from $\varphi, s, x_1, \ldots, x_i$, where φ is the content of the advice tape. The *advice complexity*, $b(n)$, of ALG is the largest number of bits of φ read by ALG over all possible inputs of length at most n.

For a request sequence σ, $\texttt{ALG}(\sigma)$ ($\texttt{OPT}(\sigma)$) denotes the non-negative *cost* of the output computed by ALG (OPT) when serving σ. We say that ALG is *c-competitive* if there exists a constant α such that $\texttt{ALG}(\sigma) \leq c \cdot \texttt{OPT}(\sigma) + \alpha$ for all request sequences σ. If the inequality holds with $\alpha = 0$, we say that ALG is *strictly c-competitive*.

A *randomized online algorithm* $\mathtt{R}$ *with advice complexity* $b(n)$ is a probability distribution over deterministic online algorithms with advice complexity at most $b(n)$. We say that $\mathtt{R}$ is c-competitive if there exists a constant α such that $\mathbb{E}[\mathtt{R}(\sigma)] \leq c \cdot \mathtt{OPT}(\sigma) + \alpha$ for all request sequences σ. △

Suppose that $\mathtt{ALG}$ is a deterministic algorithm with advice complexity $b(n)$. For every fixed input length n, $\mathtt{ALG}$ can be converted into $2^{b(n)}$ algorithms $\mathtt{ALG}_1, \ldots, \mathtt{ALG}_{2^{b(n)}}$ without advice, one for each possible advice string, such that $\mathtt{ALG}(\sigma) = \min_i \mathtt{ALG}_i(\sigma)$ for every input σ of length at most n. Another basic observation is that the advice partitions the set of inputs of length n into $2^{b(n)}$ disjoint subsets.

In order to make the definition more readable, there are some (unimportant) technicalities not addressed in Definition 1. In particular, we will assume that whenever $\mathtt{ALG}$ has to compute some y_i, there is a non-empty set of *valid* values that y_i may take (this set may depend on $s, x_1, \ldots, x_i$ and $y_1, \ldots, y_{i-1}$. As an example, on a page-fault, a paging algorithm must evict a page which is currently in its cache.). An algorithm is required to compute a valid answer in each round.

In Definition 1, the initial state, s, is part of the input and known to the online algorithm before the first request x_1 arrives. Note that the initial state is not counted as a request itself. For the k-server problem on a metric space M, the initial state is a placement of the k servers in M. For the bin packing problem, there is only one possible initial state (the empty state). For problems where there is only one possible initial state, we will usually omit it and simply write $\sigma = (x_1, \ldots, x_n)$ when specifying an input.

There are several possible definitions of a randomized algorithm with advice. The one we give in Definition 1 is as close as possible to the standard definitions used for randomized online algorithms without advice [17] and randomized online algorithms in the advice-per-request model [48]. In Appendix C.1, we discuss some of the alternative definitions and how they relate to Definition 1.

Finally, we remark that the advice is provided on an infinite tape to prevent the algorithm from learning anything from the length of the advice [26, 63]. However, all lower bounds proved in this paper also holds even if the algorithm simply receives (at the very beginning) a finite advice string of length $b(n)$ where n is the input length. See Appendix C.2 for a more detailed discussion of this point.

Parameter of the problem or part of the input? When applying our main results, Theorems 1 and 2, to a problem P, some care needs to be taken regarding exactly how P is defined. Take as an example the k-server problem. In this paper, we will not consider this problem to be one single problem, but instead a family of problems parameterized by a number, k, of servers and a finite metric space (M, d). Let $P_{k,(M,d)}$ denote the k-server problem on the metric space (M, d). A $P_{k,(M,d)}$-input $\sigma = (s_0, r_1, \ldots, r_n)$ consists of an initial placement, s_0, of the k servers and n requests. The number of possible length n inputs is $|M|^k \cdot |M|^n = |M|^{k+n}$, which is $2^{n^{O(1)}}$ since k and $|M|$ are fixed. When applying Theorem 1 to $P_{k,(M,d)}$, we are allowed to set $c = \log k$ since k is fixed and therefore $\log k$ is a constant not depending on the input σ.

Regarding the additive constant α in Definition 1, we do *not* allow α to depend on the initial state s of the input. For several problems, allowing α to depend on s would trivialize the problem (for example, this is the case for the matching problems discussed in Sections 6.5 and 6.11). We remark that several papers on the k-server problem do allow the additive constant to depend on the initial configuration. While this is important when dealing with unbounded metric spaces, it does not matter when the metric space is finite. Indeed, suppose that a k-server algorithm is c-competitive on a metric space (M, d) with an additive term $\alpha(s_0)$ depending on the initial

placement s_0 of the k servers. If the metric space is finite, there are only finitely many possible initial placements of the servers. Thus, we may remove the dependency on s_0 by replacing $\alpha(s_0)$ with the constant $\max_{s_0} \alpha(s_0)$. Generally, for any problem P with a finite number of possible initial states, our requirement that α depends only on P and not the initial state does not make any difference for the competitive ratio of P-algorithms.

2.2 Information theoretic preliminaries

We will define the quantities of information theory that we are going to use (see also [40]). In order to simplify notation, we will use the following common conventions: $0 \log 0 = 0$, $0 \log(q/0) = 0$, and $0 \log(0/q) = 0$ for any fixed $q \geq 0$.

Let Ω be a finite set. In this paper, most probability spaces will be finite (they will consist of a finite number of input sequences). A *probability distribution* over Ω is a mapping $\mu : \Omega \to [0, 1]$ such that $\sum_{\omega \in \Omega} \mu(\omega) = 1$. The *support* of μ is defined as $\text{supp}(\mu) = \{\omega \in \Omega : \mu(\omega) > 0\}$. Formally, a random variable X over a finite set Ω is a mapping $X : \Omega \to \mathcal{X}$ from Ω into some set $\mathcal{X}$. We say that $\mathcal{X}$ is the *range* of X. Random variables are always denoted by capital letters and their range by the calligraphic version of that letter. We always assume that the range[2] of a random variable X over a finite set Ω is chosen so that for every $x \in \mathcal{X}$, it holds that $\Pr[X = x] > 0$.

The *entropy*, $H(X)$, of a random variable X is defined as

$$H(X) = \sum_{x \in \mathcal{X}} \Pr[X = x] \log\left(\frac{1}{\Pr[X = x]}\right). \tag{3}$$

One important property of entropy is that $0 \leq H(X) \leq \log |\mathcal{X}|$ with equality if and only if X is uniformly distributed. The *joint entropy* $H(X_1, \ldots, X_n)$ of a sequence $X_1, \ldots, X_n$ of random variables is the entropy of the random variable $(X_1, \ldots, X_n)$. Let Y be a random variable. For any fixed $y \in \mathcal{Y}$, we define $H(X|Y = y)$ to be the entropy of the random variable X conditioned on the event $Y = y$. The *conditional entropy* $H(X|Y)$ is defined as the expected value of $H(X|Y = y)$ over all $y \in \mathcal{Y}$:

$$H(X|Y) = \mathbb{E}_y[H(X|Y = y)] = \sum_{y \in \mathcal{Y}} \Pr[Y = y] H(X|Y = y). \tag{4}$$

It holds that $H(X, Y) = H(X) + H(Y|X)$. The chain rule for conditional entropy states that

$$H(X_1, X_2, \ldots, X_n|Y) = \sum_{i=1}^{n} H(X_i|X_1, \ldots, X_{i-1}, Y). \tag{5}$$

The *mutual information* is defined as $I(X; Y) = H(X) + H(Y) - H(X, Y)$. It follows from the chain rule that $I(X; Y) = H(Y) - H(Y|X) = H(X) - H(X|Y)$. Also, it can be shown that $I(X; Y) \geq 0$.

Let $\mu, \nu : \Omega \to [0, 1]$ be two probability distributions on some finite set Ω. The *Kullback-Leibler divergence* between μ and ν is defined as

$$D_{KL}(\mu \| \nu) = \sum_{\omega \in \Omega} \mu(\omega) \log\left(\frac{\mu(\omega)}{\nu(\omega)}\right). \tag{6}$$

[2]Our definition of the range $\mathcal{X}$ of X is equivalent to what is often called the support of X.

If $\mu(\omega) = 0$ for some $\omega \in \Omega$, then the term in the sum corresponding to ω is zero according to our conventions. On the other hand, if $\nu(\omega) = 0$ and $\mu(\omega) > 0$ for some $\omega \in \Omega$, then $D(\mu\|\nu)$ is undefined. In particular, $D(\mu\|\nu)$ is only defined if $\mathrm{supp}(\mu) \subseteq \mathrm{supp}(\nu)$.

Pinsker's inequality. Let μ and ν be two probability distributions on a finite set Ω. Define the L_1-*distance* between μ and ν as

$$\|\mu - \nu\|_1 = \sum_{x \in \Omega} |\mu(x) - \nu(x)| . \tag{7}$$

Let f be a function from Ω into $\mathbb{R}$ and let $\|f\|_\infty = \max_{x \in \Omega} |f(x)|$. Then,

$$\big|\mathbb{E}_\mu[f(x)] - \mathbb{E}_\nu[f(x)]\big| = \left| \sum_{x \in \Omega} f(x)\mu(x) - \sum_{x \in \Omega} f(x)\nu(x) \right| \leq \|f\|_\infty \cdot \|\mu - \nu\|_1. \tag{8}$$

Pinsker's inequality ([40, Lemma 11.6.1])[3] says this if the Kullback-Leibler divergence between μ and ν is small, then μ and ν must be close in L_1-norm:

$$\begin{aligned} \|\mu - \nu\|_1 &\leq \sqrt{2 \ln 2 \cdot D_{\mathrm{KL}}(\mu\|\nu)} \\ &= \sqrt{\ln 4 \cdot D_{\mathrm{KL}}(\mu\|\nu)}. \end{aligned} \tag{9}$$

The function $K_y(x)$. We define a family of functions related to the Kullback-Leibler divergence.

Definition 2. For $0 < y < 1$, define $K_y : [0,1] \to \mathbb{R}$ by

$$K_y(x) = x \log\left(\frac{x}{y}\right) + (1-x) \log\left(\frac{1-x}{1-y}\right). \tag{10}$$

Note that $K_y(x)$ is the Kullback-Leibler divergence $D(\mu\|\nu)$ between μ and ν where μ (resp. ν) is the Bernoulli distribution with parameter x (resp. y).

For any fixed y, the function $K_y(x)$ is convex. Furthermore, it is decreasing for $x \in [0, y]$ and increasing for $x \in [y, 1]$. Also, $K_y(0) = \log(1/(1-y))$, $K_y(y) = 0$, and $K_y(1) = \log(1/y)$. We will often need to consider inverse functions of K_y. Unfortunately, there is no standard notation for such functions. We define $K_{y,l}^{-1}$ to be the inverse of $K_y|_{[0,y]}$ and $K_{y,r}^{-1}$ to be the inverse of $K_y|_{[y,1]}$. In order to simplify some statements, we use the following conventions: $K_{y,l}^{-1}(t) = 0$ if $t > K_y(0) = \log(1/(1-y))$ and $K_{y,r}^{-1}(t) = 1$ if $t > K_y(1) = \log(1/y)$.

Recall that for $q \in \{2, 3, \ldots\}$, the *q-ary entropy function* $h_q : [0,1] \to [0,1]$ is defined as $h_q(x) = x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x)$. A straightforward calculation reveals that $K_{1/q}(x) = (1 - h_q(1-x)) \log q$ for $q \in \mathbb{N}, q \geq 2$.

The following lemma shows that if we have an upper bound on the Kullback-Leibler divergence between μ and ν, then we can bound how much $\mu(\omega)$ can differ from $\nu(\omega)$.

[3] Some care needs to be taken when comparing different statements of Pinsker's inequality. Please note that we measure the Kullback-Leibler divergence in bits (whereas it is often measured in nats). Furthermore, the inequality is sometimes stated for a different distance measure which is half the L_1-distance.

Lemma 1. *Let $q \geq 2$ and let $\mu, \nu : \{x_1, x_2, \ldots, x_q\} \to [0,1]$ be two probability distributions such that $0 < \nu(i) < 1$ for all i. Assume that $D_{KL}(\mu \| \nu) \leq d$ for some $d \geq 0$. Then, for any $1 \leq i \leq q$, it holds that*

$$\mu(i) \geq K^{-1}_{\nu(i),l}(d), \tag{11}$$

and

$$\mu(i) \leq K^{-1}_{\nu(i),r}(d). \tag{12}$$

Proof. Fix i such that $1 \leq i \leq q$. Then,

$$\begin{aligned} D(\mu\|\nu) &= \sum_{1 \leq j \leq q} \mu(j) \log\left(\frac{\mu(j)}{\nu(j)}\right) \\ &= \mu(i)\log\left(\frac{\mu(i)}{\nu(i)}\right) + \sum_{j \neq i} \mu(j)\log\left(\frac{\mu(j)}{\nu(j)}\right) \\ &\geq \mu(i)\log\left(\frac{\mu(i)}{\nu(i)}\right) + (1-\mu(i))\log\left(\frac{1-\mu(i)}{1-\nu(i)}\right) && \{\text{Log sum inequality}\} \\ &= K_{\nu(i)}(\mu(i)). \end{aligned}$$

Thus, $K_{\nu(i)}(\mu(i)) \leq d$. We will show that (11) is true. Note that (11) is trivially true if $\nu(i) \leq \mu(i)$ since the range of $K^{-1}_{\nu(i),l}$ is $[0, \nu(i)]$. Similarly, if $d > K_{\nu(i)}(0)$, then $K^{-1}_{\nu(i),l}(d) = 0$ and so (11) is also trivially true. Assume therefore that $\mu(i) < \nu(i)$ and that $d \leq K_{\nu(i)}(0)$. In this case, it follows from $K_{\nu(i)}(\mu(i)) \leq d$ that $\mu(i) \geq K^{-1}_{\nu(i),l}(d)$ since $K_{\nu(i)}$ is decreasing on $[0, \nu(i)]$ and since $K^{-1}_{\nu(i),l}$ is the inverse of $K_{\nu(i)}$ on this interval. By a symmetric argument, the inequality (12) also holds. □

2.3 Advice complexity lower bounds via Yao's principle.

Yao's principle shows that we can prove lower bounds on randomized online algorithms with advice by proving lower bounds on deterministic algorithms with advice against a fixed input distribution. Please note that this direction of Yao's principle does not rely on the minimax theorem or other game-theoretical results (it simply relies on interchanging the order of summation).

Theorem 6 (Yao [93])**.** *Let* P *be a minimization problem, let $b(n)$ be a function and let $c > 1$. Assume that for every $\varepsilon > 0$ and every α, there is an input distribution $p_{\alpha,\varepsilon}$ such that if $\mathtt{D}$ is a deterministic* P*-algorithm with advice complexity $b(n)$ then*

$$\mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{D}(\sigma)] \geq (c-\varepsilon) \cdot \mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha. \tag{13}$$

Then, the competitive ratio of every randomized P*-algorithm with advice complexity $b(n)$ is at least c.*

Proof (following [28]). Assume by way of contradiction that there exists a randomized c'-competitive algorithm $\mathtt{R}$ with advice complexity $b(n)$ for some $c' < c$. By Definition 1, $\mathtt{R}$ defines a probability distribution μ over deterministic algorithms with advice complexity at most $b(n)$. Thus, we may write $\mathbb{E}[\mathtt{R}(\sigma)] = \mathbb{E}_{\mathtt{D}\sim\mu}[\mathtt{D}(\sigma)]$ for the expected cost incurred by $\mathtt{R}$ on a fixed input σ. By definition,

there must exist a constant α' such that $\mathbb{E}[\mathtt{R}(\sigma)] = \mathbb{E}_{\mathtt{D}\sim\mu}[\mathtt{D}(\sigma)] \leq c' \cdot \mathtt{OPT}(\sigma) + \alpha'$ for every input sequence σ.

Choose $\varepsilon = c - c'$ and $\alpha = \alpha' + 1$. Let $p_{\alpha,\varepsilon}$ be the probability distribution satisfying (13) for this choice of ε and α. Since $\mathtt{R}$ is c'-competitive, we have that $E_{\sigma\sim p_{\alpha,\varepsilon}}[\mathbb{E}_{\mathtt{D}\sim\mu}[\mathtt{D}(\sigma)]] \leq c' \, \mathbb{E}_{\sigma\sim p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha'$. On the other hand, since costs are by Definition 1 non-negative,

$$\begin{aligned}
\mathbb{E}_{\sigma\sim p_{\alpha,\varepsilon}}[\mathbb{E}_{\mathtt{D}\sim\mu}[\mathtt{D}(\sigma)]] &= \mathbb{E}_{\mathtt{D}\sim\mu}[\mathbb{E}_{\sigma\sim p_{\alpha,\varepsilon}}[\mathtt{D}(\sigma)]] && \{\text{Tonelli's theorem}\} \\
&\geq \mathbb{E}_{\mathtt{D}\sim\mu}[(c-\varepsilon)\,\mathbb{E}_{\sigma\sim p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha] && \{\text{by (13)}\} \\
&= \mathbb{E}_{\mathtt{D}\sim\mu}[c'\,\mathbb{E}_{\sigma\sim p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha' + 1] \\
&= c'\,\mathbb{E}_{\sigma\sim p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha' + 1,
\end{aligned}$$

which is a contradiction. □

3 Direct product theorems for online computation with advice

The main results of this section are two direct product theorems, namely Theorems 7 and 8. Before proving these theorems, we need to define an r-round input distribution.

3.1 Definition of r-round input distributions

Recall that informally, an r-round input distribution is an input distribution which can naturally be split up into r rounds such that in each round, the requests which are revealed are chosen independently from the requests in previous rounds. Furthermore, we are interested in the case where one can associate to each round a cost function such that the total cost of the output computed by an algorithm is simply the sum of the costs incurred in each round. We remark that a round can consist of more than one request.

Given P-inputs $\sigma_1, \ldots, \sigma_r$, we define $\sigma = \sigma_1 \ldots \sigma_r$ to be the P-input obtained by concatenating the requests of the r inputs and using the same initial state as σ_1. For example, if $\sigma_1 = (s, x_1, \ldots, x_n)$ and $\sigma_2 = (s', x'_1, \ldots, x'_{n'})$, then $\sigma_1\sigma_2 = (s, x_1, \ldots, x_n, x'_1, \ldots, x'_{n'})$. We are now ready to formally define an r-round input distribution. We will also define a number of random variables associated to such a distribution.

Definition 3 (r-round input distribution)**.** Let P be a minimization problem and let $r \in \mathbb{N}$. For each $1 \leq i \leq r$, let I_i be a finite set of P-inputs such that the following holds: If $\sigma_1, \ldots, \sigma_r$ are such that $\sigma_i \in I_i$ for $1 \leq i \leq r$, then $\sigma = \sigma_1\sigma_2 \ldots \sigma_r$ is a valid P-input. Furthermore, let $I^r = I_1 \times \cdots \times I_r = \{\sigma_1 \ldots \sigma_i \ldots \sigma_r \mid \sigma_i \in I_i \text{ for } 1 \leq i \leq r\}$.

For each $1 \leq i \leq r$, let cost_i be a function which maps an output γ computed for an input $\sigma \in I^r$ to a non-negative real number $\mathrm{cost}_i(\gamma, \sigma)$. We say that cost_i is the ith round *cost function.*

Let $p_i : I_i \to [0,1]$ be a probability distribution over I_i and let $p^r : I^r \to [0,1]$ be the (product) probability distribution which maps $\sigma_1\sigma_2 \ldots \sigma_r \in I^r$ into $p_1(\sigma_1)p_2(\sigma_2)\cdots p_r(\sigma_r)$. We say that p^r (together with the associated cost functions cost_i) is an *r-round input distribution.* For $1 \leq i \leq r$, we say that p_i is the ith round input distribution of p^r. △

Given an r-round input distribution and a fixed deterministic algorithm, we introduce a number of random variables which are useful when analyzing how the algorithm performs against the r-round input distribution. These random variables will be used throughout the paper. Recall that I^r is by Definition 3 a finite set, and so these random variables are finite.

Definition 4. Let P be a minimization problem and let $r \in \mathbb{N}$. Let $p^r : I^r \to [0,1]$ be an r-round input distribution with associated cost functions cost_i and with ith round input distributions p_i. Let ALG be a deterministic P-algorithm with advice. For $\sigma \in I^r$, let $\varphi(\sigma)$ denote the advice read by ALG on input σ. Assume that $|\varphi(\sigma)| \leq b$ for all $\sigma \in I^r$, that is, assume that ALG reads at most b bits of advice when the input belongs to I^r. We define the following random variables:

- For $1 \leq i \leq r$, let X_i be the requests[4] revealed in round i. Formally, the random variable $X_i : I^r \to I_i$ maps $\sigma = \sigma_1 \dots \sigma_r \in I^r$ to $\sigma_i \in I_i$. Furthermore, let $X = X_1 \dots X_r$.
- Let B be the advice bits read by ALG. Formally, the random variable $B : I^r \to \{0,1\}^b$ maps σ to the advice $\varphi(\sigma) \in \{0,1\}^b$ read by ALG on input σ.
- For $1 \leq i \leq r$, let $\text{cost}_i(\texttt{ALG})$ be the ith round cost function applied to the output computed by ALG. Formally, the random variable $\text{cost}_i(\texttt{ALG}) : I^r \to \mathbb{R}$ maps $\sigma \in I^r$ to $\text{cost}_i(\gamma, \sigma) \in \mathbb{R}$ where γ are the answers computed by ALG when given σ as input. Furthermore, let $\text{cost}(\texttt{ALG})$ be the total cost incurred by ALG, that is, $\text{cost}(\texttt{ALG}) = \sum_{i=1}^{r} \text{cost}_i(\texttt{ALG})$.
- Let $W_i = (X_1, \dots, X_{i-1}, B)$ be the information available to ALG when the ith round begins. This information consists of all previous requests and the advice read by ALG. Formally, $W_i : I^r \to I_1 \times \dots \times I_{i-1} \times \{0,1\}^{b(n)}$ maps $\sigma \in I^r$ to $(\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \varphi(\sigma))$.

Finally, for every $1 \leq i \leq r$ and every $w \in \mathcal{W}_i$, we define the conditional ith round input distribution $p_{i|w} : I_i \to [0,1]$ as follows[5]:

$$p_{i|w}(x) = p_i(x|W_i = w) = \frac{\Pr(X_i = x, W_i = w)}{\Pr(W_i = w)}.$$

△

Even if ALG is an algorithm without advice, we will still use the random variables introduces in Definition 4. In this case, B is the empty string and hence W_i is simply the *history* of requests revealed in previous rounds. This will be relevant when proving and applying Theorem 8.

When proving lower bounds via r-round input distributions, we always assume that the algorithm receives all of the advice at the very beginning. This only strengthens the lower bounds. However, some care has to be taken due to the fact that the length (number of requests) of an input drawn from an r-round input distribution is a random variable. By Definition 3, there is only a finite number of possible request sequences in each round and, hence, there exists an upper bound on the total length of inputs from the r-round distribution. Thus, the value b (the upper bound on the number of advice bits read by ALG) in Definition 4 is well-defined. When defining B and W_i, we assume that the algorithm ALG receives a b-bit advice string at the beginning. Please note that even if σ is such that ALG does not actually read b bits of advice when serving σ, we still assume that ALG receives the first b bits of the advice-tape prepared by its oracle for σ (again, this can only help the algorithm).

3.2 An information-theoretic direct product theorem

We prove an information theoretical direct product theorem by extending the high-entropy lower bound technique due to Emek et al. [48]. We will make use of the notation of Definitions 3 and 4.

[4]The variable X_1 will in addition to the requests revealed in round 1 also contain the initial state of the input.

[5]Recall that we are assuming that the range $\mathcal{W}_i$ of the finite random variable W_i is chosen to be smallest possible, that is, if $w \in \mathcal{W}_i$, then $\Pr[W_i = w] > 0$.

Theorem 7. *Let* P *be a minimization problem and let* p^r *be an* r*-round input distribution with associated cost functions* cost_i*. Furthermore, let* ALG *be a deterministic algorithm reading at most* b *bits of advice on every input in the support of* p^r*. Assume that there exists a convex and decreasing function* $f : [0, \infty] \to \mathbb{R}$ *such that for every* $1 \leq i \leq r$ *and* $w \in \mathcal{W}_i$*, the following holds:*

$$\mathbb{E}[\text{cost}_i(\texttt{ALG})|W_i = w] \geq f\left(D_{KL}(p_{i|w} \| p_i)\right). \tag{14}$$

Then, $\mathbb{E}[\text{cost}(\texttt{ALG})] \geq rf(b/r)$.

Before sketching the proof of Theorem 7, we informally discuss the theorem. Recall that W_i is the information available to the algorithm when round i begins (the advice read and the history of previous requests). Without any advice or knowledge of the history, the probability of $x \in \mathcal{X}_i$ being selected as the round i request sequence is $p_i(x)$. However, this probability may change given that the algorithm knows W_i (in the most extreme case, the advice could specify exactly the request sequence in round i). For any fixed $w \in \mathcal{W}_i$, the probability of x being selected in round i given that $W_i = w$ is denoted $p_{i|w}(x)$. Assumption (14) informally means that the closer $p_{i|w}$ and p_i are to each other, the better a lower bound we must have on the expected cost incurred by ALG in round i given that ALG knows w. We remark that the convexity assumption on f is automatically satisfied in most applications. The conclusion of Theorem 7 essentially says that under these assumptions, an algorithm with b bits of advice for the entire input can do no better than an algorithm with b/r bits of advice for each individual round. In particular, this will allow us to conclude that $\Omega(r)$ bits of advice are needed to get a non-trivial improvement over having no advice at all.

Proof (of Theorem 7). First of all, note that $H(B)$, the entropy of B, is at most b since the algorithm ALG reads at most b bits of advice. It follows that

$$I(X; B) = H(B) - H(B|X) \leq b. \tag{15}$$

On the other hand, using (in the third equality below) that $X_1, \ldots, X_r$ are independent random variables and the chain rule of conditional entropy, we get

$$\begin{aligned} I(X; B) &= I(X_1, X_2, \ldots, X_r; B) \\ &= H(X_1, X_2, \ldots, X_r) - H(X_1, X_2, \ldots, X_r|B) \\ &= \sum_{i=1}^{r} H(X_i) - \sum_{i=1}^{r} H(X_i|X_1, \ldots, X_{i-1}, B) \\ &= \sum_{i=1}^{r} I(X_i; B, X_1, X_2, \ldots, X_{i-1}) \\ &= \sum_{i=1}^{r} I(X_i; W_i). \end{aligned} \tag{16}$$

Combining (16) and (15) gives

$$I(X; B) = I(X_1; W_1) + I(X_2; W_2) + \cdots + I(X_r; W_r) \leq b. \tag{17}$$

Fix i such that $1 \leq i \leq r$. By definition, for any $w \in \mathcal{W}_i$,

$$p_{i|w}(x) = p_i(x|W_i = w) = \frac{\Pr(X_i = x, W_i = w)}{\Pr(W_i = w)}.$$

Note that $p_{i|w}$ is, for each fixed $w \in \mathcal{W}_i$, a probability distribution satisfying $\mathrm{supp}(p_{i|w}) \subseteq \mathrm{supp}(p_i)$. Thus, the Kullback-Leibler divergence between $p_{i|w}$ and p_i is well-defined and finite. Using the law of total expectation, the assumption (14), and Jensen's inequality, we get that

$$\begin{aligned}\mathbb{E}[\mathrm{cost}_i(\mathtt{ALG})] &= \mathbb{E}_w[\mathbb{E}[\mathrm{cost}_i(\mathtt{ALG})|W_i = w]] \\ &\geq \mathbb{E}_w\left[f\big(D_{\mathrm{KL}}(p_{i|w}\|p_i)\big)\right] \\ &\geq f\left(\mathbb{E}_w[D_{\mathrm{KL}}(p_{i|w}\|p_i)]\right). \end{aligned} \tag{18}$$

As the following calculation shows, the mutual information $I(X_i; W_i)$ is the expected Kullback-Leibler divergence between between $p_{i|w}$ and p_i over all $w \in \mathcal{W}_i$:

$$\begin{aligned} I(X_i; W_i) &= \sum_{x \in I_i,\, w \in \mathcal{W}_i} \Pr(X_i = x, W_i = w) \cdot \log\left(\frac{\Pr(X_i = x, W_i = w)}{\Pr(W_i = w)\Pr(X_i = x)}\right) \\ &= \sum_{x \in I_i,\, w \in \mathcal{W}_i} \Pr(W_i = w) \cdot p_{i|w}(x) \cdot \log\left(\frac{p_{i|w}(x)}{p_i(x)}\right) \\ &= \sum_{w \in \mathcal{W}_i} \Pr(W_i = w) \cdot D_{\mathrm{KL}}(p_{i|w}\|p_i) \\ &= \mathbb{E}_w[D_{\mathrm{KL}}(p_{i|w}\|p_i)]. \end{aligned} \tag{19}$$

Combining (18) and (19) and using linearity of expectation, we get that

$$\begin{aligned}\mathbb{E}[\mathrm{cost}(\mathtt{ALG})] &= \sum_{i=1}^{r} \mathbb{E}[\mathrm{cost}_i(\mathtt{ALG})] \\ &\geq \sum_{i=1}^{r} f\big(\mathbb{E}_w[D_{\mathrm{KL}}(p_{i|w}\|p_i)]\big) \\ &= \sum_{i=1}^{r} f(I(X_i; W_i)). \end{aligned}$$

We claim that the sum $\sum_{i=1}^{r} f(I(X_i; W_i))$ is at least $rf(b/r)$. In order to prove this claim, suppose that $h_1, \ldots, h_r$ are non-negative real numbers such that $\sum_{i=1}^{r} h_i \leq b$. Then, by Jensen's inequality and the assumption that f is convex and decreasing, it follows that

$$f\left(\frac{b}{r}\right) \leq f\left(\frac{\sum_{i=1}^{r} h_i}{r}\right) \leq \frac{1}{r}\sum_{i=1}^{r} f(h_i).$$

Thus, $rf(b/r) \leq \sum_{i=1}^{r} f(h_i)$. Since $I(X_1; W_1), \ldots, I(X_r, W_r)$ are non-negative real numbers which according to (17) sums to at most b, we conclude that $\mathbb{E}[\mathrm{cost}(\mathtt{ALG})] \geq rf(b/r)$. □

For a simple example of how to apply Theorem 7, we refer to Theorem 9.

3.3 A martingale-theoretic direct product theorem

We will now provide an alternative direct product theorem using the Azuma-Hoeffding inequality.

Theorem 8. *Let* P *be a minimization problem and let* p^r *be an* r*-round input distribution with associated cost functions* cost_i*. Assume that there exist constants* $t, s \geq 0$ *such that for every deterministic* P*-algorithm* `D` *without advice, it holds that for every* i *and every history* $w \in \mathcal{W}_i$*,*

$$\mathbb{E}[\text{cost}_i(D)|W_i = w] \geq t, \tag{20}$$

$$\mathbb{E}[(\text{cost}_i(D) - t)^2|W_i = w] \leq s^2. \tag{21}$$

Then, for every algorithm `ALG` *reading at most* b *bits of advice and every* $\varepsilon > 0$*,*

$$\Pr\left[cost(ALG) \leq (t-\varepsilon)r\right] \leq \exp_2\left(b - K_{\frac{\gamma}{1+\gamma}}\left(\frac{\alpha+\gamma}{1+\gamma}\right)r\right), \tag{22}$$

where $\gamma = \dfrac{s^2}{t^2}, \alpha = \dfrac{\varepsilon}{t}$*,* $\exp_2(x) = 2^x$*, and* $K_y(x) = x\log\left(x/y\right) + (1-x)\log\left((1-x)/(1-y)\right)$*.*

It is very important to note that (20) and (21) must only hold for an algorithm *without advice*. In this case, W_i is simply the history of requests in the previous rounds (and does not contain any advice about the input). Assumption (20) says that the expected cost of a deterministic algorithm (without advice) should be at least t in each round, no matter what has happened in previous rounds. Regarding assumption (21), note that if M is an upper bound on the cost incurred by any algorithm in a single round, then (21) is satisfied by taking $s = M$. However, in some cases, it is possible to obtain a better estimate and thereby a better advice complexity lower bound. Azuma-Hoeffding's inequality essentially shows that under these assumptions, the probability that a single fixed algorithm without advice incurs a total cost smaller than $(t-\varepsilon)r$ decreases exponentially in r. Viewing an algorithm with b bits of advice as 2^b algorithms without advice then gives the result.

Note that if $b < K_{\frac{\gamma}{1+\gamma}}\left(\frac{\alpha+\gamma}{1+\gamma}\right)r$ then the probability in (22) is strictly smaller than 1. Thus, by the probabilistic method, there exists an input σ such that the algorithm with b bits of advice incurs a cost of at least $(t-\varepsilon)r$. In particular, we again get that in order to achieve a non-trivial improvement over an algorithm without advice, it must be the case that $b = \Omega(r)$.

Proof (of Theorem 8). Let `D` be a deterministic algorithm without advice. Recall that for $1 \leq j \leq r$, the random variable $\text{cost}_j(\texttt{D})$ is the cost incurred by `D` in round j when the input is drawn from the r-round input distribution p^r. Let $V_0 = 0$ and for $1 \leq i \leq r$, let $V_i = i \cdot t - \sum_{j=1}^{i} \text{cost}_j(\texttt{D})$. We claim that $V_0, V_1, V_2, \ldots, V_r$ is a supermartingale. To this end, note that for $0 \leq i < r$,

$$\begin{aligned}\mathbb{E}[V_{i+1}|V_1, \ldots, V_i] &= \mathbb{E}[V_i + t - \text{cost}_{i+1}(\texttt{D})|V_1, \ldots, V_i] \\ &= V_i + t - \mathbb{E}[\text{cost}_{i+1}(\texttt{D})|V_1, \ldots, V_i].\end{aligned}$$

Since we assume that $\mathbb{E}[\text{cost}_{i+1}(\texttt{D})|W_{i+1} = w] \geq t$ for every $w \in \mathcal{W}_{i+1}$, it follows that the conditional expectation $\mathbb{E}[\text{cost}_{i+1}(\texttt{D})|V_1, \ldots, V_i]$ is a random variable which is always at least t. Thus, $\mathbb{E}[V_{i+1}|V_1, \ldots, V_i] \leq V_i$, which proves that $V_0, V_1, \ldots, V_r$ is a supermartingale.

Since costs are always non-negative, we have that $V_{i+1} - V_i = t - \text{cost}_{i+1}(\texttt{D}) \leq t$. Furthermore, from (21) we get that $\mathbb{E}[(V_{i+1} - V_i)^2|V_1, \ldots, V_i] = \mathbb{E}[(t - \text{cost}_{i+1}(\texttt{D}))^2|V_1, \ldots, V_i] \leq s^2$. Thus, by the

Azuma-Hoeffding inequality (see (61) in Appendix B.2),

$$\begin{aligned}\Pr[\text{cost}(\mathtt{D}) \le (t-\varepsilon)r)] &= \Pr[rt - \text{cost}(\mathtt{D}(X)) \ge \varepsilon r] \\ &= \Pr\left[V_r \ge \varepsilon r\right] \\ &\le \exp_2\left(-K_{\frac{\gamma}{1+\gamma}}\left(\frac{\alpha+\gamma}{1+\gamma}\right)r\right).\end{aligned}$$

Let $\mathtt{ALG}$ be a deterministic algorithm reading at most b bits of advice for every input in $\text{supp}(p^r)$. Recall that $\mathtt{ALG}$ can be converted into 2^b deterministic algorithms without advice. Above, we computed an upper bound on the probability that a single fixed deterministic algorithm without advice will incur a cost of at most $(t-\varepsilon)r$. Using the union bound, we can get an upper bound on the probability that at least one of the 2^b algorithms will incur at most this cost:

$$\Pr[\text{cost}(\mathtt{ALG}) \le (t-\varepsilon)r)] \le \exp_2\left(b - K_{\frac{\gamma}{1+\gamma}}\left(\frac{\alpha+\gamma}{1+\gamma}\right)r\right).$$

□

Theorem 8 can be easier to apply than Theorem 7, but it is also less flexible. In this paper, we mainly use Theorem 7, but most of the results could also be proved using Theorem 8. In some cases (such as Theorem 13), Theorem 8 would yield a slightly weaker bound than Theorem 7. Our main application of Theorem 8 is a lower bound for online graph coloring given in Section 9.

4 Lower bounds for Σ-repeatable online problems

In order to prove Theorem 1, we first define the repeated version, P^*_Σ, of an online problem P. This is a new online problem consisting of some number of (variable-length) rounds of P. Just before each round, a restart takes place and everything is reset to some initial state of P. While we are usually not interested in the problem P^*_Σ itself, we introduce it as a stepping stone for proving lower bounds for P.

Given P-inputs $\sigma_1, \ldots, \sigma_r$ with the same initial state s, we define $(\sigma_1; \ldots; \sigma_r)$ to be a sequence with the requests of the r inputs concatenated and such that the initial state s arrives together with the first request of each σ_i. For example, if $\sigma_1 = (s, x_1, \ldots, x_n)$ and $\sigma_2 = (s, x'_1, \ldots, x'_{n'})$, then $(\sigma_1; \sigma_2) = (s, \{s, x_1\}, x_2, \ldots, x_n, \{s, x'_1\}, x'_2, \ldots, x'_{n'})$. Note that if $\sigma = \sigma_1 \ldots \sigma_r$ and $\sigma^* = (\sigma_1; \ldots; \sigma_r)$, then σ is an actual P-input whereas σ^* technically is not (since $\{s, x_1\}$ is not a valid P-request). When defining the repeated version of an online problem, we need to use σ^* in order to make sure that the algorithm knows when one phase ends and another begins (this may not be possible for the algorithm to deduce in σ).

Definition 5. Let P be an online problem, let S be the set of initial states for P, and let I (I_s) be the set of all possible request sequences for P (with initial state s). Define P^*_Σ to be the online problem with inputs $I^* = \{\sigma^* = (\sigma_1; \sigma_2; \ldots; \sigma_r) \mid r \ge 1, s \in S, \sigma_i \in I_s\}$. An algorithm for P^*_Σ must produce an output $\gamma^* = (\gamma_1, \ldots, \gamma_r)$ where $\gamma_i = (y_1, \ldots, y_{n_i})$ is a valid sequence of answers for the P-input $\sigma_i = (s, x_1, \ldots, x_{n_i}) \in I_s$. The score of the output γ^* is $\text{score}(\gamma^*, \sigma^*) = \sum_{i=1}^r \text{score}_P(\gamma_i, \sigma_i)$, where $\text{score}_P(\gamma_i, \sigma_i)$ is the score of the P-output γ_i with respect to the P-input σ_i. The optimal offline algorithm for P^*_Σ is denoted $\mathtt{OPT}^*_\Sigma$.

$\mathrm{P}^*_\vee$ is defined similarly, except that $\mathrm{score}(\gamma^*, \sigma^*) = \max\{\mathrm{score}_P(\gamma_1, \sigma_1), \ldots, \mathrm{score}_P(\gamma_r, \sigma_r)\}$.

In order to better understand the definition of P^*_Σ, it is useful to imagine that after serving the last request of round $i-1$ but before serving the first request of round i, the current state is changed to the initial state s (note that the algorithm knows when this happens since in $(\sigma_1; \ldots; \sigma_r)$, the first request of each σ_i is special). It is, however, important to keep in mind that even though a reset occurs when round i begins, the previous rounds are not forgotten. The algorithm has perfect recall of all requests and answers in all previous rounds. Without this recall, advice complexity lower bounds for P^*_Σ would be easy to obtain but would be of little use.

If P is the k-server problem, then an initial state s is a placement of the k servers in the metric space. Thus, in P^*_Σ, after serving the last request of round $i-1$ and before serving the first request of round i, the k servers automatically return to their initial position specified by s. Note that when P is the k-server problem, we can concatenate P-inputs $\sigma_1, \sigma_2, \ldots, \sigma_r$ into one long P-input $\sigma = \sigma_1\sigma_2 \ldots \sigma_r$. The only difference between serving the P-input σ and serving the P^*_Σ-input $\sigma^* = (\sigma_1; \ldots; \sigma_r)$ is that for the P-input σ, the k servers are not returned to the initial state s when a round ends. However, if the underlying metric space has finite diameter Δ, this difference in the placement of servers when a new round begins can be handled by ensuring that $k\Delta$ is small compared to the total cost incurred during each round. In fact, it turns out that for many online problems, there is a natural reduction from P^*_Σ to P that essentially preserves all lower bounds. This is formalized in Definition 6.

4.1 Lower bounds for P^*_Σ

We will now show how to obtain advice complexity lower bounds for P^*_Σ by repeating a hard input distribution r times and using our information theoretical direct product theorem.

Lemma 2. *Fix $r \geq 1$. Let P be a minimization problem. Let $p : I_s \to [0,1]$ be an input distribution, where I_s is a finite set of P-inputs with the same initial state s and length at most n'. Assume that for every deterministic P-algorithm ALG without advice, it holds that $\mathbb{E}_{\sigma\sim p}[\mathtt{ALG}(\sigma)] \geq t$. Also, let M be the largest cost that any P-algorithm can incur on any input from I_s.*

Then, there exists an r-round input distribution, p^r, over P^*_Σ-inputs with at most rn' requests in total, such that any deterministic P^*_Σ-algorithm reading at most b bits of advice (on inputs of length at most rn') has an expected cost of at least $r(t - 2M\sqrt{b/r})$. Furthermore, $\mathbb{E}_{\sigma^*\sim p^r}[\mathtt{OPT}^*_\Sigma(\sigma^*)] = r\,\mathbb{E}_{\sigma\sim p}[\mathtt{OPT}(\sigma)]$.

Proof. Define $p^r : \{(\sigma_1; \ldots; \sigma_r) \mid \sigma_i \in I_s\} \to [0,1]$ to be the P^*_Σ-input distribution which maps $\sigma^* = (\sigma_1; \sigma_2; \ldots; \sigma_r)$ into $p(\sigma_1)p(\sigma_2)\cdots p(\sigma_r)$. Thus, in each of the r rounds, we draw independently a request sequence from I_s according to p. From the definition of P^*_Σ, we naturally obtain a cost-function, cost_i, for each round $1 \leq i \leq r$. Together with these cost-functions, the input distribution p^r is an r-round input distribution for P^*_Σ. Note that the ith round input distribution p_i is simply $p_i = p$, and that a round of p^r corresponds to a round of P^*_Σ.

Let $\mathtt{ALG}^*$ be a deterministic algorithm for P^*_Σ reading at most b bits of advice on inputs of length at most rn'. Fix $1 \leq i \leq r$ and $w \in \mathcal{W}_i$. Let $d = D_{\mathrm{KL}}(p_{i|w} \| p_i)$. In order to apply Theorem 1, we need a lower bound on $\mathbb{E}[\mathrm{cost}_i(\mathtt{ALG}^*)|W_i = w]$ in terms of d. By Pinsker's inequality (9),

$$\|p_{i|w} - p_i\|_1 \leq \sqrt{d \cdot \ln 4}. \tag{23}$$

Thus, it suffices to bound $\mathbb{E}[\text{cost}_i(\texttt{ALG}^*)|W_i = w]$ in terms of the L_1-distance between $p_{i|w}$ and p_i. Let $h(d) = \sqrt{d \cdot \ln 4}$. We construct a P-algorithm, $\texttt{ALG}_w$, without advice by hard-wiring w (i.e., the advice and the requests in rounds 1 to $i-1$) into the P^*_Σ-algorithm $\texttt{ALG}^*$. That is, for an input sequence $\sigma \in \text{supp}(p_{i|w})$, the new algorithm $\texttt{ALG}_w$ simulates the computation of $\texttt{ALG}^*$ on σ when $W_i = w$ and $\texttt{ALG}^*$ is given the requests σ in round i. Note that this is possible since w is fixed, and hence the output of $\texttt{ALG}^*$ in round i given σ as input in this round is completely determined. For all other input sequences, $\texttt{ALG}_w$ behaves arbitrarily (but does compute some valid output). It follows that $\texttt{ALG}_w$ is well-defined for all input sequences in $\text{supp}(p_i)$. Thus, $\texttt{ALG}_w$ defines a mapping $\sigma \mapsto \texttt{ALG}_w(\sigma)$ on $\text{supp}(p_i)$ such that $\|\texttt{ALG}_w\|_\infty = \max_{\sigma \in \text{supp}(p_i)} |\texttt{ALG}_w(\sigma)| \leq M$ and such that if $\sigma \in \text{supp}(p_{i|w}) \subseteq \text{supp}(p_i)$, then $\texttt{ALG}_w(\sigma)$ is equal to the cost incurred by $\texttt{ALG}^*$ if $W_i = w$ and σ is given as input in round i. It follows that

$$\begin{aligned}
\mathbb{E}[\text{cost}_i(\texttt{ALG}^*)|W_i = w] &= \mathbb{E}_{\sigma \sim p_{i|w}}[\texttt{ALG}_w(\sigma)] \\
&\geq \mathbb{E}_{\sigma \sim p}[\texttt{ALG}_w(\sigma)] - M \cdot h(d) && \{\text{By (23) and (8)}\} \\
&\geq t - M \cdot h(d).
\end{aligned}$$

Define $f(d) = t - M \cdot \sqrt{\ln 4 \cdot d}$. Since f is convex and decreasing, it follows from our direct product theorem (Theorem 7) that (remember that $p_i = p_{\alpha,\varepsilon}$):

$$\mathbb{E}_{\sigma^* \sim p^r}[\texttt{ALG}^*(\sigma^*)] = \mathbb{E}[\text{cost}(\texttt{ALG}^*)] \geq r f(b/r) = r\left(t - M\sqrt{\frac{b \cdot \ln 4}{r}}\right) \geq r\left(t - 2M\sqrt{\frac{b}{r}}\right). \tag{24}$$

Here, we used that $\sqrt{\ln 4} < 2$. Finally, we observe that

$$\begin{aligned}
\mathbb{E}_{\sigma^* \sim p^r}[\texttt{OPT}^*_\Sigma(\sigma^*)] = \mathbb{E}_{\sigma^* \sim p^r}\left[\sum_{i=1}^r \texttt{OPT}(\sigma_i)\right] &= \sum_{i=1}^r \mathbb{E}_{\sigma^* \sim p^r}[\texttt{OPT}(\sigma_i)] = \sum_{i=1}^r \mathbb{E}_{\sigma_i \sim p_i}[\texttt{OPT}(\sigma_i)] \\
&= \sum_{i=1}^r \mathbb{E}_{\sigma_i \sim p}[\texttt{OPT}(\sigma_i)] = r\,\mathbb{E}_{\sigma \sim p}[\texttt{OPT}(\sigma)].
\end{aligned}$$

□

4.2 Definition of repeatable online problems

Lemma 2 shows how to obtain lower bounds for P^*_Σ. Of course, we are usually not very interested in the problem P^*_Σ itself. However, for several online problems, there is an obvious reduction from P^*_Σ to P since it is possible for an adversary to simulate a restart in P. This idea is formalized by the notion of a repeatable online problem.

Definition 6. Let P be an online minimization problem such that for every fixed P-input, there is only a finite number of valid outputs. Furthermore, let $k_1, k_2, k_3 \geq 0$. We say that P is Σ-*repeatable with parameters* (k_1, k_2, k_3) if there exists a mapping $g : I^* \to I$ with the following properties:

Σ1. For every $\sigma^* \in I^*$,

$$|g(\sigma^*)| \leq |\sigma^*| + k_1 r, \tag{25}$$

where r is the number of rounds in σ^*.

Σ2. For every deterministic P-algorithm $\texttt{ALG}$, there is a deterministic P^*_Σ-algorithm $\texttt{ALG}^*$ such that for every $\sigma^* \in I^*$,

$$\texttt{ALG}^*(\sigma^*) \leq \texttt{ALG}(g(\sigma^*)) + k_2 r, \tag{26}$$

where r is the number of rounds in σ^*.

Σ3. For every $\sigma^* \in I^*$,

$$\texttt{OPT}^*_\Sigma(\sigma^*) \geq \texttt{OPT}(g(\sigma^*)) - k_3 r, \tag{27}$$

where r is the number of rounds in σ^*.

If a problem P is repeatable with $k_2 = k_3 = 0$, we say that P is *strictly repeatable*. The definition of $\vee$-repeatable is identical to that of Σ-repeatable, except that P^*_Σ and $\texttt{OPT}^*_\Sigma$ are replaced by $\mathrm{P}^*_\vee$ and $\texttt{OPT}^*_\vee$. The following simple lemma is useful when working with repeatable online problems.

Lemma 3. *Let* P *be a* Σ*-repeatable problem with parameters* (k_1, k_2, k_3)*, let* b *be a constant, and let* $\texttt{ALG}$ *be a* P*-algorithm with advice. Let* $J^* \subseteq I^*$ *be a subset of* P^*_Σ*-inputs. Suppose that for every* $\sigma \in J^*$*, the algorithm* $\texttt{ALG}$ *reads at most* b *bits of advice on the* P*-input* $g(\sigma^*)$*. Then there exists a* P^*_Σ*-algorithm* $\texttt{ALG}^*$ *such that* $\texttt{ALG}^*$ *reads exactly* b *bits of advice on every input* $\sigma^* \in J^*$ *and such that* $\texttt{ALG}(g(\sigma^*)) \geq \texttt{ALG}^*(\sigma^*) - k_2 r$ *for every input* $\sigma^* \in J^*$ *where* r *is the number of rounds in* σ^**.*

Proof. By definition, there exists 2^b deterministic algorithms, $\texttt{ALG}_1, \ldots, \texttt{ALG}_{2^b}$, without advice such that $\texttt{ALG}(g(\sigma^*)) = \min_i \texttt{ALG}_i(g(\sigma^*))$ for every input $\sigma^* \in J^*$. For each $1 \leq i \leq 2^b$, it follows from (26) in Definition 6 that there exists a deterministic P^*_Σ-algorithm $\texttt{ALG}^*_i$ without advice such that $\texttt{ALG}_i(g(\sigma^*)) \geq \texttt{ALG}^*_i(\sigma^*) - k_2 r$ for every σ^*. We are now ready to define $\texttt{ALG}^*$ and the corresponding advice oracle: The algorithm $\texttt{ALG}^*$ always reads b bits of advice before the first request arrives. This is possible since b is a constant. For an input $\sigma^* \in J^*$, the oracle computes an index i such that $\texttt{ALG}(g(\sigma^*)) = \texttt{ALG}_i(g(\sigma^*))$. This index i is written onto the advice tape. The algorithm $\texttt{ALG}^*$ learns i from the advice tape and serves σ^* using $\texttt{ALG}^*_i$. It follows that $\texttt{ALG}(g(\sigma^*)) \geq \texttt{ALG}^*(\sigma^*) - k_2 r$ for any input $\sigma^* \in J^*$. □

We remark that Lemma 3 will only be used when the length of inputs in J^* is bounded. This explains why the lemma is only stated for the case where b is a constant. Also, note that for inputs outside of J^*, the algorithm $\texttt{ALG}^*$ can behave arbitrarily.

4.3 Proof of Theorem 1

We are now ready to prove our main theorem for Σ-repeatable online problems, Theorem 1. For convenience, we restate the theorem before giving the proof.

Theorem 1. *Let* P *be a* Σ*-repeatable online minimization problem and let* c *be a constant. Suppose that for every* $\varepsilon > 0$ *and every* α*, there exists an input distribution* $p_{\alpha,\varepsilon} : I \to [0,1]$ *with finite support such that* $\mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{D}(\sigma)] \geq (c - \varepsilon)\,\mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{OPT}(\sigma)] + \alpha$ *for every deterministic algorithm* $\texttt{D}$ *without advice. Then, every randomized algorithm reading at most* $o(n)$ *bits of advice on inputs of length* n *has a competitive ratio of at least* c*.*

Proof. Let P be a Σ-repeatable online problem with parameters (k_1, k_2, k_3). The lower bound will be obtained via Yao's principle (Theorem 6). To this end, fix $\varepsilon' > 0$ and α' and let $\texttt{ALG}$ be an arbitrary deterministic P-algorithm reading $b(n) \in o(n)$ bits of advice. We need to show that there exists a probability distribution p such that $\mathbb{E}_p[\texttt{ALG}(\sigma)] \geq (c - \varepsilon')\,\mathbb{E}_p[\texttt{OPT}(\sigma)] + \alpha'$. Choose $\varepsilon = \varepsilon'$

and $\alpha = \alpha' + k_2 + k_3(c-\varepsilon) + 1$. The reason for choosing this particular value of α will become clear later on. For this choice of ε and α, let $p_{\alpha,\varepsilon}$ be a probability distribution for which it holds that $\mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{D}(\sigma)] \geq (c-\varepsilon)\,\mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha$ for every deterministic algorithm $\mathtt{D}$ without advice. Without loss of generality, we assume that all inputs in the support of $p_{\alpha,\varepsilon}$ have the same initial state (Lemma 17 in Appendix E justifies this assumption).

P is Σ-repeatable, and so for any fixed input in $\mathrm{supp}(p_{\alpha,\varepsilon})$ there is only a finite number of possible outputs. Since $\mathrm{supp}(p_{\alpha,\varepsilon})$ is itself finite, this means that there exists a real number M such that, on inputs from $\mathrm{supp}(p_{\alpha,\varepsilon})$, no valid P-algorithm incurs a cost larger than M[6]. Thus, from $p_{\alpha,\varepsilon}$ we obtain for each $r \in \mathbb{N}$ via Lemma 2 a probability distribution $p^r_{\alpha,\varepsilon}$ such that for any P^*_Σ-algorithm $\mathtt{ALG}^*$ which on inputs in $\mathrm{supp}(p^r_{\alpha,\varepsilon})$ reads at most $b^* \in \mathbb{N}$ bits of advice, it holds that $\mathbb{E}_{\sigma^* \sim p^r_{\alpha,\varepsilon}}[\mathtt{ALG}^*(\sigma^*)] \geq r\left((c-\varepsilon)\,\mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha - 2M\sqrt{b^*/r}\right)$.

Since $\mathrm{supp}(p_{\alpha,\varepsilon})$ is finite, there exists a constant $L_{\alpha,\varepsilon}$ such that every $\sigma \in \mathrm{supp}(p_{\alpha,\varepsilon})$ contains at most $L_{\alpha,\varepsilon}$ requests. It follows that the number of requests in any input sequence in the support of $p^r_{\alpha,\varepsilon}$ is at most $L_{\alpha,\varepsilon} r$. By (25) in Definition 6, this implies that the number of requests in $g(\sigma^*)$ is at most $L_{\alpha,\varepsilon} r + k_1 r$ for every $\sigma^* \in \mathrm{supp}(p^r_{\alpha,\varepsilon})$. Forget for a moment that b is the advice complexity of $\mathtt{ALG}$, and just view b as a function $b : \mathbb{N} \to \mathbb{N}$. By assumption, $b(n) \in o(n)$. Since k_1 and $L_{\alpha,\varepsilon}$ are constants (independent of r), this implies that $b(L_{\alpha,\varepsilon} r + k_1 r)/r \to 0$ as $r \to \infty$. Choose r large enough that $2M\sqrt{\frac{b(L_{\alpha,\varepsilon} r + k_1 r)}{r}} \leq 1$.

We have now fixed the value of r. Let $b_r = b(L_{\alpha,\varepsilon} r + k_1 r)$. Note that b_r is a fixed integer (since r is fixed) and that the algorithm $\mathtt{ALG}$ will read at most b_r bits of advice on any input $g(\sigma^*)$ where $\sigma^* \in \mathrm{supp}(p^r_{\alpha,\varepsilon})$. Using Lemma 3, we convert the P-algorithm $\mathtt{ALG}$ into a P^*_Σ-algorithm $\mathtt{ALG}^*$ such that $\mathtt{ALG}^*$ reads b_r bits of advice on every input $\sigma^* \in \mathrm{supp}(p^r_{\alpha,\varepsilon})$ and such that $\mathtt{ALG}(g(\sigma^*)) \geq \mathtt{ALG}^*(\sigma^*) - k_2 r$ for every $\sigma^* \in \mathrm{supp}(p^r_{\alpha,\varepsilon})$. The proof is completed by the following calculation:

$$
\begin{aligned}
\mathbb{E}_{\sigma^* \sim p^r_{\alpha,\varepsilon}}[\mathtt{ALG}(g(\sigma^*))] &\geq \mathbb{E}_{\sigma^* \sim p^r_{\alpha,\varepsilon}}[\mathtt{ALG}^*(\sigma^*)] - k_2 r \\
&\geq r\left((c-\varepsilon)\,\mathbb{E}_{\sigma \sim p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha - 2M\sqrt{\frac{b_r}{r}}\right) - k_2 r && \{\text{Lemma 2}\} \\
&\geq r\Big((c-\varepsilon)\,\mathbb{E}_{\sigma \sim p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha - 1\Big) - k_2 r && \{\text{by choice of } r\} \\
&= (c-\varepsilon) r\,\mathbb{E}_{\sigma \sim p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] - r(1 + k_2 - \alpha) \\
&= (c-\varepsilon)\,\mathbb{E}_{\sigma^* \sim p^r_{\alpha,\varepsilon}}[\mathtt{OPT}^*_\Sigma(\sigma^*)] - r(1 + k_2 - \alpha) && \{\text{Lemma 2}\} \\
&\geq (c-\varepsilon)\,\mathbb{E}_{\sigma^* \sim p^r_{\alpha,\varepsilon}}[\mathtt{OPT}(g(\sigma^*))] - r(1 + k_2 + k_3(c-\varepsilon) - \alpha) && \{\text{by (27)}\} \\
&\geq (c-\varepsilon')\,\mathbb{E}_{\sigma^* \sim p^r_{\alpha,\varepsilon}}[\mathtt{OPT}(g(\sigma^*))] + \alpha'. && \{\text{by choice of } \alpha\}
\end{aligned}
$$

□

[6]Here, we use that according to Definition 1, the cost of an output cannot be ∞ but must be a real number.

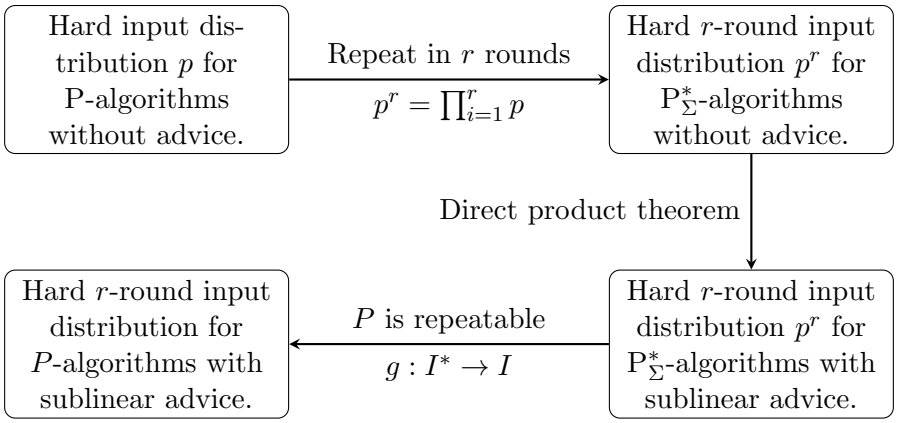
Hard input distribution p for P-algorithms without advice.
Repeat in r rounds
$p^r = \prod_{i=1}^{r} p$
Hard r-round input distribution p^r for P^*_Σ-algorithms without advice.
Direct product theorem
Hard r-round input distribution for P-algorithms with sublinear advice.
P is repeatable
$g : I^* \to I$
Hard r-round input distribution p^r for P^*_Σ-algorithms with sublinear advice.

FIGURE 1: Schematic overview of our technique for obtaining advice complexity lower bounds for repeatable online problems. One starts with a hard input distribution p for P-algorithms without advice, then repeats it in r rounds to get a hard r-round input distribution p^r for P^*_Σ. Using a direct product theorem, we get that p^r is also a hard input distribution for algorithms with sublinear advice. These two steps are collected in Lemma 2. Finally, since P is repeatable, we can obtain a lower bound on P-algorithms with sublinear advice using the reduction g.

4.4 Generalized task systems are repeatable

We seek a simple set of conditions which are sufficient to show that an online problem is Σ-repeatable. It is not hard to show that any problem which can be modeled as a task system [28] is Σ-repeatable. However, it turns out that one can relax the definition of a task system and still maintain repeatability. Therefore, we introduce *generalized task systems (GTS)*. Essentially, a GTS is a task system with (almost) no restrictions imposed on the distance function. Lemma 4 shows that every online problem which can be modeled as a GTS satisfying a certain finiteness condition is Σ-repeatable.

Definition 7. A *generalized task system (GTS)* is a minimization problem defined by a triple $(\mathcal{S}, \mathcal{T}, d)$ where $\mathcal{S}$ is a set of N *states*, $\mathcal{T}$ is a set of allowable tasks and $d : \mathcal{S} \times \mathcal{S} \to \mathbb{R}_{\geq 0}$ is an arbitrary function. A *task* is a mapping $t : \mathcal{S} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ satisfying that there exists at least one state $s \in \mathcal{S}$ such that $t(s) \neq \infty$.

An input $\sigma = (s_0, t_1, t_2, \ldots, t_n)$ consists of an *initial state* $s_0 \in \mathcal{S}$ and n tasks, $t_1, \ldots, t_n$, where $t_i \in \mathcal{T}$ for each $1 \leq i \leq n$. An online algorithm must compute as output a set of states $y = (s_1, \ldots, s_n)$ such that s_i is computed from $(s_0, t_1, t_2, \ldots, t_i)$ and such that $t_i(s_i) \neq \infty$. When the algorithm outputs a state s_i in round i, it incurs a cost of $d(s_{i-1}, s_i) + t_i(s_i)$. We say that $d(s_{i-1}, s_i)$ is the *transition cost* and that $t_i(s_i)$ is the *processing cost* of serving the task t_i. The total cost of the output y is the sum of the costs incurred in each round.

If the function d is clear from the context, we will sometimes write $(\mathcal{S}, \mathcal{T})$ instead of $(\mathcal{S}, \mathcal{T}, d)$. If for a task t and a state s it holds that $d(s, s) + t(s) = 0$, then we say that s is a *haven* against the task t. Note that if $d(s, s) > 0$, then s cannot be a haven against any task.

Clearly, any task system and any metrical task system (MTS) is also a generalized task system (GTS). Recall that for any MTS with N states, there is a $O(\text{polylog } N)$-competitive randomized algorithm [50]. It is easy to see that no such result can be achieved for an arbitrary GTS. However,

if an online problem P can be modeled as a GTS, this gives us a lot of information about the (game-theoretical) structure of P: It gives us a notion of a state, it shows that the cost of an output can be decomposed into the costs of each answer, and other similar properties. It will often be important for us to have upper (and lower) bounds on the cost that can be incurred while serving a single task.

Definition 8. Let $(\mathcal{S}, \mathcal{T}, d)$ be a GTS. The *max-cost* of $(\mathcal{S}, \mathcal{T}, d)$ is the infimum over all real numbers Δ which for every $t \in \mathcal{T}$ satisfy the following condition:

$$\text{For all } s, s' \in \mathcal{S}, \text{ if } t(s') \neq \infty \text{ then } d(s, s') + t(s') \leq \Delta. \tag{28}$$

The *min-cost* of $(\mathcal{S}, \mathcal{T}, d)$ is the supremum over all real numbers δ which for every $t \in \mathcal{T}$ satisfy the following conditions:

$$\text{For all } s, s' \in \mathcal{S}, \text{ if } s \neq s' \text{ then } d(s, s') \geq \delta. \tag{29}$$

$$\text{For all } s \in \mathcal{S}, \text{ either } d(s, s) + t(s) = 0 \text{ or } d(s, s) + t(s) \geq \delta. \tag{30}$$

If $(\mathcal{S}, \mathcal{T})$ has max-cost Δ then it follows from (28) that the cost incurred by any algorithm in a single round is at most Δ. Similarly, if $(\mathcal{S}, \mathcal{T})$ has min-cost δ, then (29) ensures that if an online algorithm switches to a different state, then it incurs a cost of at least δ. The condition (30) ensures that if an online algorithm stays in the same state in some round, then it either incurs no cost at all or it incurs a cost of at least δ. The min-cost will play an important role later on in Section 5. For now, we only need to consider generalized task systems with bounded max-cost.

Lemma 4. *Let* P *be a minimization problem which can be modeled as a generalized task system with finite max-cost* Δ*. Then* P *is* Σ*-repeatable with parameters* $(0, \Delta, \Delta)$.

Proof. Let P be the online problem defined by the generalized task system $(\mathcal{S}, \mathcal{T}, d)$. Assume that P has finite max-cost Δ. We will show that P is Σ-repeatable with parameters $(0, \Delta, \Delta)$. Let I (resp. I^*) be the set of all input sequences for P (resp. P*). An input $\sigma^* = (\sigma_1; \ldots; \sigma_r) \in I^*$ consists of r rounds. For $1 \leq i \leq r$, $\sigma_i = (s_0, t_1^i, t_2^i, \ldots, t_{n_i}^i)$ is a sequence of tasks belonging to $\mathcal{T}$.

The mapping $g : I^* \to I$ maps $\sigma^* = (\sigma_1; \sigma_2; \ldots; \sigma_r)$ into the request sequence

$$\sigma = \sigma_1 \sigma_2 \cdots \sigma_r = \left(s_0, t_1^1, t_2^1, \ldots, t_{n_1}^1, t_1^2, \ldots, t_{n_2}^2, \ldots, t_1^r, \ldots, t_{n_r}^r\right). \tag{31}$$

Note that σ is obtained from σ^* simply by revealing the same tasks in the same order but not performing the reset to the initial state s_0 at the beginning of each round $i > 1$. It is clear that (25) of Definition 6 is satisfied with $k_1 = 0$ since $|g(\sigma^*)| = |\sigma^*|$. Let ALG be a deterministic P-algorithm (without advice). Define ALG* to be the following P*-algorithm: ALG* serves the jth task of σ^* by moving to the same state as ALG does when serving the jth task of $g(\sigma^*)$. We will now prove that so defined ALG* satisfies (26). When a reset occurs at the end of a round, ALG* will be moved from its current state to the initial state of the next round (at no cost). This means that when serving the very first task after a reset, ALG* might incur a larger cost than ALG does. However, by assumption, the cost incurred by ALG* for serving the first request of the next round is at most Δ. For the remaining tasks in the round, ALG* incurs exactly the same cost as ALG. Thus, $\texttt{ALG}^*(\sigma^*) \leq \texttt{ALG}(g(\sigma^*)) + k_2 r$ where $k_2 = \Delta$ is a constant independent of ALG* and r.

The proof of (27) is similar. OPT can serve the tasks of $g(\sigma^*)$ in the same way as $\texttt{OPT}_\Sigma^*$ serves the corresponding tasks of σ^*. The only problem is that when a reset occurs, $\texttt{OPT}_\Sigma^*$ is moved to the initial state of the new round for free, which might allow $\texttt{OPT}_\Sigma^*$ to serve the very first task after the reset at a lower cost than OPT. However, this difference in cost for serving the first request of a round can be at most Δ. Thus, $\texttt{OPT}_\Sigma^*(\sigma^*) \geq \texttt{OPT}(g(\sigma^*)) - k_3 r$ where $k_3 = \Delta$. □

5 Compact online problems

It is natural to ask when a lower bound on randomized online algorithm without advice is witnessed by a family of input distributions satisfying the conditions of Theorem 1 (we formalize this question in Definition 9). It is well-known that there exist problems with an infinite number of requests and answers in each round for which an optimal lower bounds on randomized algorithms cannot be obtained via Yao's principle (see e.g. [27]). However (as we will see in Section 5.2), for several important online problems, one can show that a lower bound on randomized algorithms must always be witnessed by a family of input distributions which are compatible with Theorem 1.

5.1 Definition of compact online problems and proof of Theorem 2

Definition 9. Let P be a minimization problem and let $c > 1$ be a constant such that the expected competitive ratio of every randomized P-algorithm is at least c. We say that P is *compact* if for every $\varepsilon > 0$ and every $\alpha \geq 0$, there exists an input distribution $p_{\alpha,\varepsilon}$ with finite support such that if $\mathtt{D}$ is a deterministic online algorithm (without advice), then $\mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{D}(\sigma)] \geq (c-\varepsilon) \cdot \mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha$.

Informally, an online problem is compact if we can use inputs of bounded length to prove lower bounds on the competitive ratio which are essentially tight. The following two lemmas and Theorem 2 follows trivially by combining Theorem 1 with Definition 9. Nevertheless, we believe that Theorem 2 is an important conceptual implication of Theorem 1.

Lemma 5. *Let* P *be a compact and* Σ*-repeatable online problem, and let* $c > 1$ *be a constant. Assume that every randomized algorithm without advice has a competitive ratio of at least* c*. Then the competitive ratio of every randomized algorithm reading* $o(n)$ *bits of advice is at least* c*.*

Proof. The problem P is compact, and so for every $\varepsilon > 0$ and every $\alpha \geq 0$, there exists a probability distribution $p_{\alpha,\varepsilon}$ with finite support such that $\mathbb{E}_{p_{\alpha,\varepsilon}}[\mathtt{D}(\sigma)] \geq (c-\varepsilon) E_{p_{\alpha,\varepsilon}}[\mathtt{OPT}(\sigma)] + \alpha$ for every deterministic algorithm $\mathtt{D}$. Since P is Σ-repeatable, it follows from Theorem 1 that every randomized algorithm with sublinear advice complexity has a competitive ratio of at least c. □

Lemma 6. *Let* P *be a compact and* Σ*-repeatable online problem and let* $c \geq 1$ *be a constant. If there exists a (possibly randomized)* c*-competitive algorithm reading* $o(n)$ *bits of advice then, for all* $\varepsilon > 0$*, there exists a randomized* $(c+\varepsilon)$*-competitive algorithm without advice.*

Proof. Let $\mathtt{ALG}$ be a c-competitive P-algorithm reading $o(n)$ bits of advice. Assume by way of contradiction that for some $\varepsilon > 0$, there does not exist a randomized $(c+\varepsilon)$-competitive algorithm without advice. By Lemma 5, this implies that the competitive ratio of $\mathtt{ALG}$ must be at least $c+\varepsilon$, a contradiction. □

Theorem 2. *Let* P *be a compact and* Σ*-repeatable minimization problem with at most* $2^{n^{O(1)}}$ *inputs of length* n*, and let* c *be a constant not depending on* n*. The following are equivalent:*

1. *For every* $\varepsilon > 0$*, there exists a randomized* $(c+\varepsilon)$*-competitive* P*-algorithm without advice.*
2. *For every* $\varepsilon > 0$*, there exists a deterministic* $(c+\varepsilon)$*-competitive* P*-algorithm with advice complexity* $o(n)$*.*

Proof. Suppose that for every $\varepsilon > 0$, there exists a randomized $(c+\varepsilon)$-competitive P-algorithm without advice. Since P is a minimization problem with at most $2^{n^{O(1)}}$ possible inputs of length n, it follows from the derandomization result of Böckenhauer et al. [25] (or Theorem 16) that for every $\varepsilon > 0$, there exists a $(c+2\varepsilon)$-competitive deterministic P-algorithm with advice complexity $O(\log n)$. Since ε was arbitrary, this proves the first implication of the theorem.

Suppose that for every $\varepsilon > 0$, there exists a deterministic $(c+\varepsilon)$-competitive P-algorithm with advice complexity $o(n)$. Then, by Lemma 6, for every $\varepsilon > 0$ there exists a randomized $(c+2\varepsilon)$-competitive P-algorithm without advice. Since ε was arbitrary, this proves the remaining implication of the theorem. □

Theorem 2 is of course only interesting in those cases where one can prove that a problem is compact without actually determining the best possible competitive ratio that a randomized algorithm can achieve. We now sketch how to prove that several important Σ-repeatable online problems are compact. Interestingly, our proof will rely on the "upper bound part" of Yao's principle [93] which is (much) less frequently used than the lower bound part.

Fix a Σ-repeatable problem P such that for every n, the number of inputs of length at most n is finite[7]. Let $c > 1$ be a constant such that the expected competitive ratio of every randomized P-algorithm is at least c. What does it mean for P to *not* be compact? It means that there exists an $\varepsilon > 0$ and $\alpha \geq 0$ such that the following holds: For every $n' \in \mathbb{N}$ and for every probability distribution p over P-inputs of length at most n', there exists a deterministic algorithm $\mathtt{D}$ such that $\mathbb{E}_{\sigma\sim p}[\mathtt{D}(\sigma)] < (c-\varepsilon)\,\mathbb{E}_{\sigma\sim p}[\mathtt{OPT}(\sigma)] + \alpha$. Recall that, by assumption, there is only a finite number of inputs and outputs for P of length at most n'. This makes it possible to view the problem P restricted to inputs of length at most n' as a *finite* two-player zero-sum game between an algorithm and adversary. Thus (see Lemma 10 for the full proof), by Yao's principle, we get that there exists a randomized P-algorithm $\mathtt{R}_{n'}$ such that $\mathbb{E}[\mathtt{R}_{n'}(\sigma)] < (c-\varepsilon)\mathtt{OPT}(\sigma) + \alpha$ for every P-input of length at most n'. Now, if we can somehow show that it is possible to use the algorithms $\mathtt{R}_1, \mathtt{R}_2, \ldots$ to obtain a single algorithm $\mathtt{R}$ which is better than c-competitive (on all possible inputs), then the problem at hand must, by contradiction, be compact. In what follows, we will use this strategy to show that all problems which can be modeled as a task system satisfying certain requirements are compact.

5.2 Lazy task systems are compact

Recall that we introduced generalized task systems in Definition 7 in order to make it easier to show that an online problem is Σ-repeatable. We will now define a certain subclass of generalized task systems known as *lazy task systems (LTS)*. In a lazy task system, the distance function must satisfy the triangle inequality and it must separate states. Theorem 6 shows that these conditions on the distance function together with the finiteness of the state space imply that an online problem is compact (and Σ-repeatable). In Appendix D, we show how to model several well-known online problems as lazy task systems. We note that a metrical task system is always a lazy task system.

Definition 10. A *lazy task system (LTS)* is a generalized task system $(\mathcal{S}, \mathcal{T}, d)$ where the function

[7]Since P is Σ-repeatable (Definition 6), this assumption automatically implies that there is only finitely many different P-algorithms for inputs of length at most n.

d satisfies the following conditions for all $s, s', s'' \in \mathcal{S}$:

$$d(s, s') > 0 \text{ if } s \neq s'. \tag{32}$$
$$d(s, s') \leq d(s, s'') + d(s'', s'). \tag{33}$$

We say that a LTS-algorithm $\texttt{ALG}$ is *lazy* if $\texttt{ALG}$ does not move to another state if its current state is a haven for the current task.

The reason for the name *lazy* task system is that given a LTS-algorithm $\texttt{ALG}$, there exists a lazy LTS-algorithm $\texttt{ALG}'$ such that $\texttt{ALG}'(\sigma) \leq \texttt{ALG}(\sigma)$ for every input σ. This follows since the function d satisfies the triangle inequality (33). We will need this fact when proving that a finite LTS is compact.

The following lemma states that a finite LTS always has a strictly positive min-cost (and a bounded max-cost). The lemma follows immediately from the finiteness of the LTS combined with inequality (32).

Lemma 7. *Let $(\mathcal{S}, \mathcal{T})$ be a LTS where $\mathcal{S}$ and $\mathcal{T}$ are both finite sets. There exists $\delta, \Delta \in \mathbb{R}$ such that $0 < \delta \leq \Delta$ and such that $(\mathcal{S}, \mathcal{T})$ has min-cost δ and max-cost Δ.*

We now turn to proving that finite lazy task systems are compact. As previously mentioned, we will use the same proof technique as Ambühl [5] and Mömke [84]. We begin by introducing the necessary terminology.

Definition 11. Let $x \in \mathbb{N}$ and let A be a set of deterministic algorithms $\{\texttt{ALG}_1, \texttt{ALG}_2, \ldots\}$ for a LTS $(\mathcal{S}, \mathcal{T})$. Let $\sigma = (s_0, t_1, \ldots, t_n)$ be a $(\mathcal{S}, \mathcal{T})$-input. An algorithm is *caught* after serving t_i if it has incurred a cost of at least x for serving $t_1, \ldots, t_i$. After serving a request, an algorithm which has not yet been caught can choose to *surrender*. An algorithm which is not caught nor has surrendered is said to be *free*. A request is *bad* (with respect to A and x) if at least one free algorithm of A incurs a non-zero cost for serving the request.

We remark that eventually we will only consider algorithms that never surrenders. However, allowing an algorithm to surrender makes it easier to state a sufficiently strong induction hypothesis in the proof of the following lemma.

Lemma 8. *For every $x \in \mathbb{R}_{\geq 0}$ and $k \in \mathbb{N}$, there exists some $N_{k,x} \in \mathbb{N}$ such that for every finite LTS $(\mathcal{S}, \mathcal{T})$ with min-cost δ and $|\mathcal{S}| = k$, and every set $A = \{\texttt{ALG}_1, \texttt{ALG}_2, \ldots\}$ of lazy deterministic algorithms (each of which is initially placed in some state), no input sequence can contain more than $\delta^{-k} N_{k,x}$ bad requests with respect to A and x.*

Proof. Since the LTS $(\mathcal{S}, \mathcal{T})$ is finite, its min-cost δ is strictly positive according to Lemma 7. We start by assuming that $\delta = 1$.

Fix $x \in \mathbb{R}_{\geq 0}$. The proof is by induction on k. Setting $N_{1,x} = \lceil x \rceil$ shows that the lemma is true for $k = 1$. Indeed, since there is one unique state, all free algorithms incurs a cost of at least 1 for each bad task (since we are assuming that $\delta = 1$). Thus, after $N_{1,x} = \lceil x \rceil$ bad tasks, all algorithms which have not surrendered will have incurred a cost of at least $\lceil x \rceil \geq x$.

Assume the lemma is true for some $k \in \mathbb{N}$, and let $(\mathcal{S}, \mathcal{T})$ be a LTS with $|\mathcal{S}| = k + 1$. Fix a set of algorithms A. Let σ be an input sequence and partition σ into phases $\sigma_1, \sigma_2, \ldots, \sigma_l$ (depending on A) as follows: The first phase σ_1 starts with the first task of σ. A phase ends when there have

been $N_{k,x}+1$ bad tasks since the beginning of the phase. Thus, each complete phase contains exactly $N_{k,x}+1$ bad tasks (the last task is always bad) and an arbitrary number of good tasks. Note that $\sigma = \sigma_1\sigma_2\ldots\sigma_l$, where the final phase σ_l may be incomplete. Let $1 \leq j < l$. Assume that $\sigma_j = (t_1, \ldots, t_m)$. We want to show that the following claim is true.

Claim: If an algorithm in the set A is free at the beginning of a complete phase σ_j, then that algorithm must either surrender or incur a cost of at least 1 while serving σ_j.

Note that the last task t_m of the phase σ_j must be bad since σ_j is a complete phase. Let $\tau = (t_1, \ldots, t_{m-1})$ be all tasks of σ_j except the final bad task t_m. In order to prove the claim, assume that $\texttt{ALG}' \in A$ is an algorithm which does not incur any cost for serving τ and which is still free after serving t_{m-1} (if no such algorithm exists, the claim is trivially true). We want to show that $\texttt{ALG}'$ must necessarily incur a cost of 1 for serving t_m. Note that $\texttt{ALG}'$ must have been in the same state, s', while serving all of the tasks in τ (because of inequality (32)) and that s' must have been a haven (Definition 7) for all tasks in τ. In particular, $d(s', s') = 0$. In order to show that $\texttt{ALG}'$ incurs a cost of at least 1 for serving t_m, we will show that all free algorithms in A must be in state s' after serving t_{m-1}. If there exists a task $t \in \tau$ such that $t(s) = \infty$ for all $s \neq s'$, then this is obviously true. Suppose therefore that for every task $t \in \tau$, at least one state in $\mathcal{S} \setminus \{s'\}$ is available.

Let $M \subseteq A$ be the set of free algorithms which were not in state s' at the beginning of the phase σ_j. Let $M_{s'} \subseteq M$ be those algorithms in M which at some point during the phase served a task in τ by moving to s', and therefore (since all algorithms in A are lazy) stayed in s' while serving τ. We will now make use of the induction hypothesis. Suppose that we modify all algorithms in M such that if $\texttt{ALG} \in M$ moves to s' when serving some task, the modified algorithm, $\widetilde{\texttt{ALG}}$, instead moves to some arbitrary available state in $\mathcal{S} \setminus \{s'\}$ and then surrenders. Let $\widetilde{M}$ be the set of modified algorithms. The algorithms in $\widetilde{M}$ are well-defined and valid algorithms for the LTS $(\mathcal{S} \setminus \{s'\}, \widetilde{\mathcal{T}})$, where $\widetilde{\mathcal{T}}$ consists of the same tasks as $\mathcal{T}$ but with the state s' removed. Consider now the task sequence $\widetilde{\tau}$ for $(\mathcal{S} \setminus \{s'\}, \widetilde{\mathcal{T}})$ induced by τ. If $t_i \in \tau$ is bad with respect to A, this is because t_i is bad for some free algorithm $\texttt{ALG} \in M$ which has not yet moved into state s' while serving τ. Thus, the corresponding task $\widetilde{t_i}$ in $\widetilde{\tau}$ must be bad for the corresponding (free) algorithm $\widetilde{\texttt{ALG}}$ in $\widetilde{M}$. It follows that $\widetilde{\tau}$ must contain at least $N_{k,x}$ tasks which are bad with respect to $\widetilde{M}$. By the induction hypothesis, after the $N_{k,x}$ bad tasks of $\widetilde{\tau}$ have been served, all of the modified algorithms in $\widetilde{M}$ have either surrendered or have incurred a cost of at least x and therefore been caught. This means that for the $(\mathcal{S}, \mathcal{T})$-algorithms in M, after the first $N_{k,x}$ bad tasks of τ, all of the algorithms in M have either moved into state s', surrendered, or been caught. Recall that by definition, all algorithms in $A \setminus M$ have been in state s' since the beginning of the current phase. We conclude that when t_m arrives, all free algorithms must be in state s'. Therefore, in order for the final task t_m of the phase to be bad, that task must force all algorithms in state s' to incur a cost of at least 1. Since all other algorithms have either surrendered or been caught, this means that all algorithms which are free after serving σ_j must have incurred a cost of at least 1 during the phase σ_j. This proves the claim.

It follows that every input sequence contains at most $\lceil x \rceil$ complete phases (since after this number of complete phases, all algorithms in A have incurred a cost of at least x or have surrendered). Hence, an input sequence can contain at most $N_{k+1,x} := \lceil x \rceil \cdot (N_{k,x}+1)$ bad tasks with respect to A and x. Since this recursively defined bound on $N_{k,x}$ depends only on k and x (and not the set of algorithms A nor the LTS $(\mathcal{S}, \mathcal{T})$), this finishes the proof for $\delta = 1$. The general case follows by multiplying by δ^{-1} in each of the k steps of the induction (and hence gives a factor of δ^{-k}). $\square$

We are now ready to prove an important *Reset Lemma* for lazy task systems: Suppose that for every n', we have a good algorithm for inputs of length at most n'. Then, we can split the input into epochs of bounded length. When a new epoch begins, the good online algorithm wipes its own memory and serves the epoch as if the input had only consisted of this epoch. This results in a good algorithm for inputs of arbitrary lengths. While this sound similar to the concept of a repeatable online problem, note that for repeatable online problems, it is the *adversary* that can simulate a reset of the game. On the other hand, for LTS, it is the algorithm that can decide to essentially reset the game between itself and the adversary. The only difficulty in proving the reset lemma lies in the fact that an input sequence could contain arbitrarily long subsequences for which $\texttt{OPT}$ pays nothing. However, using Lemma 8, we can show that there is a limit as to how long the adversary can benefit from such subsequences.

Lemma 9 (Reset Lemma)**.** *Let $(\mathcal{S}, \mathcal{T})$ be a LTS where both $\mathcal{S}$ and $\mathcal{T}$ are finite. Let $c \geq 1$ be a constant and let $\alpha : \mathbb{N} \to \mathbb{R}_{\geq 0}$ be a function of n such that $\alpha = o(n)$. Assume that for every $n' \in \mathbb{N}$, there exists a randomized algorithm $\texttt{R}_{n'}$ (depending on n') such that $\mathbb{E}[\texttt{R}_{n'}(\sigma)] \leq c\texttt{OPT}(\sigma) + \alpha(n')$ for every input sequence σ where $|\sigma| \leq n'$. Then for every constant $\varepsilon > 0$, there exists a single randomized algorithm $\texttt{R}$ such that $\texttt{R}$ is $(c+\varepsilon)$-competitive (for inputs of arbitrary length).*

In this section (where our goal is to prove Theorem 2), we will only use the Reset Lemma when α is just an additive constant, that is, $\alpha = O(1)$. However, we prove a more general version of the Reset Lemma where the additive term α is allowed to be any function of order $o(n)$. This will be needed in Section 7 where we use online learning to give a constructive version of Theorem 2. Note that the lemma is obviously false if α was allowed to be of order $O(n)$ instead of $o(n)$, since no algorithm would ever incur a cost larger than $\alpha(n) = \Delta \cdot n = O(n)$ on inputs of length at most n, where Δ is the max-cost of the LTS.

Proof (of Lemma 9). Let $\varepsilon > 0$. The goal is to design an algorithm $\texttt{R}$ for inputs σ of arbitrary length. To this end, we first fix some constants (i.e., numbers which may depends on c, α, ε, and $(\mathcal{S}, \mathcal{T})$, but not on the input σ). Let $k = |\mathcal{S}|$, and let δ be the min-cost and Δ the max-cost of $(\mathcal{S}, \mathcal{T})$. By Lemma 7, we get that $0 < \delta \leq \Delta < \infty$. Furthermore, for reasons that will become clear later on, let $x = \varepsilon^{-1} k \Delta$ and choose $n' \in \mathbb{N}$ large enough so that

$$n' \geq \lceil \varepsilon^{-1}\delta^{-1}(\alpha(n') + (c+\varepsilon)\Delta) \rceil \cdot (\delta^{-k} N_{k,x} + k). \tag{34}$$

This is possible since, by assumption, $\alpha(n) \in o(n)$ while $\varepsilon, c, \delta, k$, and $N_{k,x}$ are all constants with respect to the input length n. Thus, $\lceil \varepsilon^{-1}\delta^{-1}(\alpha(n)+(c+\varepsilon)\Delta) \rceil \cdot (\delta^{-k} N_{k,x}+k) \in o(n)$, from which the existence of an n' satisfying (34) follows. Please note that n' is simply some fixed integer, depending on the constants $\varepsilon, c, \delta, \Delta, k, N_{k,x}$ and the function α. Thus, n' and $\alpha(n')$ are both independent of the input σ and its length.

By assumption, there exists a randomized algorithm $\texttt{R}_{n'}$ such that $\mathbb{E}[\texttt{R}_{n'}(\sigma)] \leq c \cdot \texttt{OPT}(\sigma) + \alpha(n')$ for every input σ of length $|\sigma| \leq n'$. Recall that formally, a randomized algorithm is a probability distribution over a (possibly infinite number of) deterministic algorithms. However, since the number of possible inputs and outputs of length at most n' is finite (we assumed that $\mathcal{S}$ and $\mathcal{T}$ were both finite), it follows that there only exist a finite number of possible deterministic algorithms for inputs of length at most n'. In particular, there must be a finite number of deterministic algorithms $\texttt{ALG}_1, \ldots, \texttt{ALG}_m$ such that for inputs of length at most n', the algorithm $\texttt{R}_{n'}$ can be viewed as a probability distribution over these m deterministic algorithms. For $1 \leq i \leq m$, let $p(i)$

be the probability assigned to the deterministic algorithm $\texttt{ALG}_i$ by $\texttt{R}_{n'}$. We will assume that the algorithms $\texttt{ALG}_1, \ldots, \texttt{ALG}_m$ are lazy. If not, we can simply make them lazy and thereby obtain a new randomized algorithm with exactly the same performance guarantee as R.

Consider an input sequence σ of arbitrary length n. We partition σ into phases as follows: The first phase begins with the first request. A phase ends when for every state $s \in \mathcal{S}$, it holds that s was not a haven for at least one task in the current phase. At any time step, we let M be the set of states which have so far been havens against every task in the current phase. Note that both the phases and the set M can be computed online. Also, note that immediately before a phase starts, every state belongs to M. During a phase, the number of states in M will decrease. When the last task of the phase arrives, M will become empty. Note that any algorithm (including OPT) must incur a cost of at least the min-cost δ while serving the tasks of a complete phase.

We will now modify the algorithm $\texttt{R}_{n'}$. The modified algorithm, $\texttt{R}'$, will still only be used on inputs of length at most n' but will serve as an important subroutine in the final algorithm for inputs of arbitrary length. In order to define $\texttt{R}'$, we first obtain a set of deterministic algorithms $\{\texttt{ALG}'_1, \ldots, \texttt{ALG}'_m\}$ by modifying the algorithms $\{\texttt{ALG}_1, \ldots, \texttt{ALG}_m\}$. For each i, the modified algorithm, $\texttt{ALG}'_i$, is defined as follows: At the beginning of each phase, $\texttt{ALG}'_i$ behaves exactly as $\texttt{ALG}_i$ (recall that an online algorithm can figure out itself when an old phase ends and a new phase begins). However, assume that a task t arrives such that after serving t, the algorithm $\texttt{ALG}_i$ will have incurred a total cost of at least x in the current phase. In this case, $\texttt{ALG}'_i$ serves t in the same way that $\texttt{ALG}_i$ does and then goes *OPT-chasing*. An algorithm which is OPT-chasing will move to some (arbitrarily chosen) state s currently in M. The algorithm will remain in state s until a task against which s is not a haven arrives. When this happens, the OPT-chasing algorithm $\texttt{ALG}'_i$ moves to another state currently in M. This continues until M is empty, that is, until the phase ends. When the last task t of the phase arrives, the algorithm $\texttt{ALG}'_i$ will stop its OPT-chasing and will serve t by moving to the same state as $\texttt{ALG}_i$ does when serving t.

Let $\widehat{\sigma}$ be a phase of σ. We claim that $\texttt{ALG}'_i(\widehat{\sigma}) \leq (1+\varepsilon)\texttt{ALG}_i(\widehat{\sigma})$. Let $\widehat{\sigma} = \widehat{\sigma}_1\widehat{\sigma}_2$ where $\widehat{\sigma}_1$ are the tasks served by $\texttt{ALG}'_i$ before it went OPT-chasing and $\widehat{\sigma}_2$ are the tasks served by $\texttt{ALG}'_i$ while OPT-chasing. By definition, $\texttt{ALG}'_i(\widehat{\sigma}_1) = \texttt{ALG}_i(\widehat{\sigma}_1)$. Also, $\texttt{ALG}'_i(\widehat{\sigma}_2) \leq k\Delta$. Since $\texttt{ALG}'_i$ only goes OPT-chasing if $\texttt{ALG}_i(\widehat{\sigma}_1) \geq x = \varepsilon^{-1}k\Delta$, we get that $\texttt{ALG}'_i(\widehat{\sigma}) \leq \texttt{ALG}_i(\widehat{\sigma}_1) + k\Delta \leq (1+\varepsilon)\texttt{ALG}_i(\widehat{\sigma}_1) \leq (1+\varepsilon)\texttt{ALG}_i(\widehat{\sigma})$.

We are now ready to define $\texttt{R}'$. The algorithm $\texttt{R}'$ randomly selects an algorithm, $\texttt{ALG}'$, from $\{\texttt{ALG}'_1, \ldots, \texttt{ALG}'_m\}$ such that $\Pr[\texttt{ALG}' = \texttt{ALG}'_i] = p(i)$. Note that $\texttt{ALG}'$ is a random variable. Also, $\texttt{R}'$ simulates (deterministically) all of the algorithms $\texttt{ALG}'_1, \ldots, \texttt{ALG}'_m$ while serving the input. If a task t arrives for which at least one of the m algorithms will incur a non-zero cost (i.e., t is bad), then $\texttt{R}'$ passes on the task t to $\texttt{ALG}'$ and serves t by moving to the same state as $\texttt{ALG}'$ does. $\texttt{R}'$ also updates its deterministic simulation of the m algorithms. On the other hand, if a task t arrives for which none of the m algorithms would incur a non-zero cost (i.e., every algorithm is in a state which is a haven against t), then the task t is *ignored*. That is, $\texttt{R}'$ does not pass on the task t to $\texttt{ALG}'$ and it does not update its deterministic simulation of the m algorithms. Instead, $\texttt{R}'$ simply serves t without incurring any cost (recall that $\texttt{R}'$ must currently be in a state which is a haven against t) and without the help of $\texttt{ALG}'$. Thus, all further requests will be handled exactly the same as if t had never arrived.

We will now bound the number of requests that can be passed on to $\texttt{ALG}'$ in a single phase. By Lemma 8, after $\texttt{R}'$ has passed on $\delta^{-k}N_{k,x}$ requests in a single phase to $\texttt{ALG}'$, all of the m algorithms $\texttt{ALG}'_1, \ldots, \texttt{ALG}'_m$ must have incurred a cost of at least x. Thus, they must have gone OPT-chasing. In particular, they are all currently occupying some state that belongs to M. Thus, from now on,

every request not ignored by $\mathtt{R}'$ will decrease the size of M by one. It follows that after passing on at most k further requests, M is empty and the phase has ended. We conclude that at most $\delta^{-k}N_{k,x} + k$ requests in a single phase can be passed on to $\mathtt{ALG}'$.

Finally, we are ready to define the algorithm $\mathtt{R}$. To this end, we divide the input sequence into *epochs*. Each epoch will consist of a number of phases. The first epoch starts when the first task of σ arrives. An epoch ends when it contains $\lceil \varepsilon^{-1}\delta^{-1}(\alpha(n') + (c+\varepsilon)\Delta)\rceil$ complete phases. Since the beginning and end of each phase can be computed online, so can the beginning and end of each epoch.

Suppose a new epoch τ has just begun. The algorithm $\mathtt{R}$ will be in some state s (the state in which $\mathtt{R}$ served the last task of the previous epoch, or, if τ is the very first epoch, s will be the initial state of the input sequence σ). The algorithm $\mathtt{R}$ will use a *new instantiation* of $\mathtt{R}'$ (which will initially be in state s) to serve τ. By a new instantiation, we mean that an epoch should be served using an instantiation of $\mathtt{R}'$ which is independent (and unaware) of all previous epochs of σ and how they were served. Formally, $\mathtt{R}$ will simulate $\mathtt{R}'$ on the $(\mathcal{S}, \mathcal{T})$-input with initial state s and tasks τ, and $\mathtt{R}$ will serve all tasks of τ exactly as they are served by $\mathtt{R}'$. When τ ends, $\mathtt{R}$ throws away the current instantiation of $\mathtt{R}'$ and serves the next epoch using a new instantiation of $\mathtt{R}'$.

We denote by $\mathtt{OPT}(\tau)$ the cost incurred by $\mathtt{OPT}$ while serving an epoch τ when given σ as input. Note that the expected cost incurred by $\mathtt{R}$ while serving an epoch τ, which we denote by $\mathbb{E}[\mathtt{R}(\tau)]$, is exactly $\mathbb{E}[\mathtt{R}'(\tau)]$. We will give an upper bound on $\mathbb{E}[\mathtt{R}'(\tau)]$ in terms of $\mathtt{OPT}(\tau)$. Since at most $N_{k,x} + k$ tasks of each phase are passed on to $\mathtt{ALG}'$ by $\mathtt{R}'$, this means that if we use $\mathtt{R}'$ to serve the epoch τ, then at most $\lceil \varepsilon^{-1}\delta^{-1}(\alpha(n') + (c+\varepsilon)\Delta)\rceil \cdot (\delta^{-k}N_{k,x} + k)$ tasks are being passed on to $\mathtt{ALG}'$ by the definition of an epoch. By the choice of n', this means that at most n' tasks are being passed on to $\mathtt{ALG}'$ in a single epoch. Note that by the triangle inequality, the optimal cost of serving an epoch must be at least the optimal cost of serving those requests in the epoch that were passed on to $\mathtt{ALG}'$. Thus, by the performance guarantee of $\mathtt{R}_{n'}$ for inputs of length at most n', we get that $\mathbb{E}[\mathtt{R}'(\tau)] \leq (c+\varepsilon)\mathtt{OPT}_s(\tau) + \alpha(n')$ where $\mathtt{OPT}_s$ is an optimal offline algorithm which initially is in the same state s as $\mathtt{R}'$ was when τ began. With the exception of the very first epoch, it could be the case that $\mathtt{OPT}$ (the optimal offline algorithm for the entire sequence σ) begins the epoch τ in a state different from s. However, we have that $\mathtt{OPT}_s(\tau) \leq \mathtt{OPT}(\tau) + \Delta$. Thus,

$$\mathbb{E}[\mathtt{R}(\tau)] = \mathbb{E}[\mathtt{R}'(\tau)] \leq (c+\varepsilon)\mathtt{OPT}_s(\tau) + \alpha(n') \leq (c+\varepsilon)\mathtt{OPT}(\tau) + \alpha(n') + (c+\varepsilon)\Delta. \tag{35}$$

Let $\tau_1, \ldots, \tau_l$ be all of the epochs (so that $\sigma = \tau_1 \ldots \tau_l$). For $1 \leq i < l$, τ_i is a complete epoch. Since $\mathtt{OPT}$ incurs a cost of at least δ for each complete phase, it follows that

$$\begin{aligned} \mathtt{OPT}(\tau_i) &\geq \delta\lceil \varepsilon^{-1}\delta^{-1}(\alpha(n') + (c+\varepsilon)\Delta)\rceil \\ &\geq \varepsilon^{-1}(\alpha(n') + (c+\varepsilon)\Delta). \end{aligned} \tag{36}$$

Combining (35) and (36), we get that $\mathtt{R}(\tau_i) \leq (c+2\varepsilon)\mathtt{OPT}(\tau_i)$ for $1 \leq i < l$. Note that the last epoch, τ_l, may be incomplete and so we cannot use this estimate for $\mathtt{R}(\tau_l)$. However, from the bound on the complete epochs and the competitiveness of $\mathtt{R}'$, it follows that

$$\begin{aligned}
\mathbb{E}[\mathtt{R}(\sigma)] &= \sum_{i=1}^{l} \mathbb{E}[\mathtt{R}(\tau_i)] = \left(\sum_{i=1}^{l-1} \mathbb{E}[\mathtt{R}(\tau_i)]\right) + \mathbb{E}[\mathtt{R}(\tau_l)] \\
&\leq \left(\sum_{i=1}^{l-1} (c+2\varepsilon)\mathtt{OPT}(\tau_i)\right) + (c+\varepsilon)\mathtt{OPT}(\tau_l) + \alpha(n') + (c+\varepsilon)\Delta \\
&\leq \left(\sum_{i=1}^{l} (c+2\varepsilon)\mathtt{OPT}(\tau_i)\right) + \alpha(n') + (c+\varepsilon)\Delta \\
&= (c+2\varepsilon)\mathtt{OPT}(\sigma) + \alpha(n') + (c+\varepsilon)\Delta.
\end{aligned}$$

Since $\alpha(n') + (c+\varepsilon)\Delta = O(1)$ does not depend on the input σ, this shows that $\mathtt{R}$ is $(c+2\varepsilon)$-competitive for input sequences of arbitrary length. Since ε was arbitrary this proves the theorem. □

We say that an online problem is *truly finite* if there are only a finite number of possible inputs and outputs. Note that we do not just require that for some fixed input length there should be a finite number of inputs (and outputs), or that there are only finitely many types of requests. It is the total number of inputs and outputs which must be finite. In particular, a necessary condition for an online problem to be finite is that there is a fixed upper bound on the number of requests. Also, all valid algorithms for a finite online problem are (non-strictly) 1-competitive (simply choose the additive constant to be sufficiently large). The following lemma follows directly from (the difficult direction of) Yao's principle. For completeness, we include the proof.

Lemma 10. *Let* P *be a truly finite minimization problem and let $c > 1$, $\varepsilon > 0$ and $\alpha \geq 0$. Assume that, for every probability distribution p over inputs to* P*, there exists a deterministic algorithm ALG such that $\mathbb{E}_{\sigma\sim p}[\mathit{ALG}(\sigma)] < (c-\varepsilon)\,\mathbb{E}_{\sigma\sim p}[\mathit{OPT}(\sigma)] + \alpha$. Then, there exists a randomized algorithm $\mathtt{R}$ such that $\mathbb{E}[\mathtt{R}(\sigma)] < (c-\varepsilon)\mathit{OPT}(\sigma) + \alpha$ for every input sequence σ.*

Proof. Let $\sigma_1, \sigma_2, \ldots, \sigma_m$ be an enumeration of all P-inputs and let $\mathtt{ALG}_1, \mathtt{ALG}_2, \ldots, \mathtt{ALG}_{m'}$ be all of the valid deterministic P-algorithms (there are only a finite number of such inputs and algorithms since P is truly finite). It is well-known that competitive analysis of an online problem may be viewed as a two-player zero-sum game between an online player and an adversary. Consider the function $f(i,j) = (c-\varepsilon)\mathtt{OPT}(\sigma_j) + \alpha - \mathtt{ALG}_i(\sigma_j)$. The payoff function for the online player is $f(i,j)$ while the payoff function for the adversary is $-f(i,j)$. Thus, the online player is trying to select i so as to maximize $f(i,j)$ while the adversary is trying to select j so as to minimize $f(i,j)$.

Since we are interested in randomized algorithms, we allow the players to use mixed strategies (a probability distribution over deterministic strategies). We denote by $\mathbf{y} = (y_1, \ldots, y_{m'})$ a mixed strategy of the online player (where y_i is the probability assigned to $\mathtt{ALG}_i$) and by $\mathbf{x} = (x_1, \ldots, x_m)$ a mixed strategy of the adversary. Let $F(\mathbf{y}, \mathbf{x}) = \mathbb{E}_{i\sim\mathbf{y}}[\mathbb{E}_{j\sim\mathbf{x}}[f(i,j)]]$ be the expected payoff (for the online player). By a slight abuse of notation, we will write $F(\mathtt{ALG}_i, \mathbf{x})$ as a shorthand for $F(\mathbf{y}, \mathbf{x})$ where $\mathbf{y}$ is the probability distribution which assigns a probability of 1 to $\mathtt{ALG}_i$. Similarly for $F(\mathbf{y}, \sigma_j)$.

Since P is a finite two-person zero-sum game, it follows from von Neumann's minimax theorem that

$$\max_{\mathbf{y}} \min_{\mathbf{x}} F(\mathbf{y}, \mathbf{x}) = \min_{\mathbf{x}} \max_{\mathbf{y}} F(\mathbf{y}, \mathbf{x}). \tag{37}$$

As a corollary (sometimes known as Loomis' lemma), it follows that there exists an optimal mixed strategy, $\mathbf{y}^\star$, for the online player and an optimal mixed strategy, $\mathbf{x}^\star$, for the adversary such that

$$\max_i F(\texttt{ALG}_i, \mathbf{x}^\star) = \min_j F(\mathbf{y}^\star, \sigma_j). \tag{38}$$

We claim that $\max_i F(\texttt{ALG}_i, \mathbf{x}^\star) > 0$. If not, then $\max_i F(\texttt{ALG}_i, \mathbf{x}^\star) \leq 0$ which implies that $\mathbb{E}_{j\sim\mathbf{x}^\star}[\texttt{ALG}_i(\sigma_j)] \geq (c-\varepsilon)\,\mathbb{E}_{j\sim\mathbf{x}^\star}[\texttt{OPT}(\sigma_j)] + \alpha$ for every deterministic algorithm $\texttt{ALG}_i$, but this contradicts the assumption of the theorem. This shows that the claim is true and thus, by (38), we have that $\min_j F(\mathbf{y}^\star, \sigma_j) = \max_i F(\texttt{ALG}_i, \mathbf{x}^\star) > 0$. Equivalently, $\mathbb{E}_{i\sim\mathbf{y}^\star}[\texttt{ALG}_i(\sigma)] < (c-\varepsilon)\texttt{OPT}(\sigma) + \alpha$ for every input σ. Thus, the mixed strategy $\mathbf{y}^\star$ proves the existence of a randomized algorithm with the desired performance guarantee.

□

Theorem 6. *Let $(\mathcal{S}, \mathcal{T})$ be a LTS such that both $\mathcal{S}$ and $\mathcal{T}$ are finite. Then $(\mathcal{S}, \mathcal{T})$ is compact.*

Proof. Assume by way of contradiction that $(\mathcal{S}, \mathcal{T})$ is a finite task system which is not compact. Let $c > 1$ be a constant such that the expected competitive ratio of every randomized $(\mathcal{S}, \mathcal{T})$-algorithm is at least c (if no such c exists, then we are done). Since $(\mathcal{S}, \mathcal{T})$ is not compact, there exists an $\varepsilon > 0$ and an $\alpha \geq 0$ such that the following holds: For every $n' \in \mathbb{N}$ and every probability distribution p over inputs of length at most n', there exists a deterministic algorithm $\texttt{ALG}$ such that

$$\mathbb{E}_p[\texttt{ALG}(\sigma)] < (c-\varepsilon)\,\mathbb{E}_p[\texttt{OPT}(\sigma)] + \alpha. \tag{39}$$

Recall that $\mathcal{S}$ and $\mathcal{T}$ are both finite by assumption. This means that for every $n' \in \mathbb{N}$, we obtain a truly finite online problem by considering the set of $(\mathcal{S}, \mathcal{T})$-inputs of length at most n'. Thus, it follows from Lemma 10 and (39) that for every $n' \in \mathbb{N}$, there exists a randomized algorithm $\texttt{R}_{n'}$ such that $\mathbb{E}[\texttt{R}_{n'}(\sigma)] < (c-\varepsilon)\texttt{OPT}(\sigma) + \alpha$ for every $(\mathcal{S}, \mathcal{T})$-input σ of length at most n'. But then Lemma 9 implies that there exists a single randomized $(\mathcal{S}, \mathcal{T})$-algorithm which is $(c-\varepsilon/2)$-competitive for inputs of arbitrary length. This gives the desired contradiction, since the competitive ratio of every randomized $(\mathcal{S}, \mathcal{T})$-algorithm was assumed to be at least c. □

6 Applications of main results

We are finally ready to apply Theorems 1 and 2. In particular, we will prove the lower bounds for Σ-repeatable problems stated in Table 1 and, when relevant, discuss the equivalence of randomization and sublinear advice implied by Theorem 2.

If a problem is Σ-repeatable but not known to be compact, we have to verify that the lower bound on randomized algorithms without advice is compatible with Theorem 1. This is often straightforward, even if the lower bound is not explicitly proved using Yao's principle. Indeed, most lower bound proofs starts by fixing (an upper bound on) the number of requests n. For this fixed n, a *finite* number of request sequences of length at most n are constructed in an adversarial manner. It is then shown that the lower bound follows by letting n tend to infinity. By Yao's principle (see Lemma 10), such a lower bound can always be converted into a family of input distributions compatible with Theorem 1.

6.1 Paging

The paging problem with a cache of size k (and a set of N pages) can be modeled as a finite LTS and is therefore compact and Σ-repeatable. Thus, the well-known (see [53, 28]) lower bound of $H_k = \Omega(\log k)$ on the competitive ratio of randomized online algorithms without advice carries over to online algorithms with sublinear advice. Previously, the best known lower bound for algorithms with sublinear advice was 5/4 [26].

It is shown in [26] that with $\lceil \log k \rceil$ bits of advice, a deterministic algorithm can achieve a competitive ratio of $3 \log k + O(1) = O(\log k)$, and with n bits of advice, a deterministic algorithm can be strictly 1-competitive. Combining these results with our lower bound, we see that the asymptotic advice complexity of paging is now fully understood: $O(1)$ bits of advice are enough for a deterministic algorithm to be $O(\log k)$-competitive (a randomized algorithm can be $O(\log k)$-competitive without advice). However, $o(n)$ bits of advice are not enough to be better than $\Omega(\log k)$-competitive, while n bits of advice suffice for a deterministic algorithm to be strictly 1-competitive.

Theorem 12 gives a lower bound on the exact number of advice bits needed to achieve some competitive ratio $1 \leq c < H_k$ for paging.

6.2 k-server

For every k and every finite metric space (M, d), the k-server problem on (M, d) can be modeled as a finite LTS and is therefore compact and Σ-repeatable. Since paging corresponds to the k-server problem on the uniform metric, the lower bound of $H_k = \Omega(\log k)$ for paging algorithms with advice complexity $o(n)$ also applies to the k-server problem. Previously, the best known lower bound for the competitive ratio of k-server algorithms with sublinear advice was 3/2 [90]. This lower bound was shown to hold even for two servers on a line, and improved an earlier lower bound of 5/4 obtained in [26] and [59].

For the 2-server problem on the line, a lower bound of $e/(e-1) \approx 1.58$ for randomized algorithms without advice is well-known [67]. For the general 2-server problem, an even better lower bound of $1 + e^{-1/2} \approx 1.60$ for randomized algorithms was obtained in [39]. Both of these lower bounds are obtained using finite metric spaces. Thus, combining these results with Theorem 1 improves on the previous best lower bound of 3/2 for the 2-server problem (on a line and in general).

In [15], Bartal et al. show that the competitive ratio of a randomized k-server algorithm on *any* metric space (with more than k points) is at least $\Omega(\log k / \log^2 \log k)$. By Theorem 1, the same lower bound holds for algorithms with advice complexity $o(n)$. No advice complexity lower bounds for the k-server problem on arbitrary metric spaces were previously known.

For arbitrary metric spaces, the best known k-server algorithm is the deterministic work-function algorithm which is $(2k-1)$-competitive [75]. In a recent breakthrough, Bansal et al. gave a randomized $O(\log^2 k \log^3 m \log \log m)$-competitive k-server algorithm [13], where m is the number of points in the underlying metric space. In particular, this algorithm has a better competitive ratio than the work-function algorithm if m is subexponential in k. It has been conjectured that it is in fact possible to achieve a competitive ratio of $O(\log k)$ for any metric space (see e.g. [74]). It is known that it is possible to achieve a competitive ratio of $O(\log k)$ using $2n$ bits of advice [25] (see also [87]). By Theorem 2, achieving the same competitive ratio with $o(n)$ bits of advice would confirm the randomized k-server conjecture (for finite metric spaces).

6.3 Metrical task systems

Any finite MTS with N states can, obviously, be modeled as a finite LTS and is therefore compact and Σ-repeatable. As with the k-server problem, the best known lower bound for randomized algorithms is the $\Omega(\log N)$ lower bound on the uniform metric space. Together with Theorem 1, this implies that a MTS-algorithm with advice complexity $o(n)$ must have a competitive ratio of at least $\Omega(\log N)$. However, this advice complexity lower bound was already obtained (for another metric space) in [48] by a reduction from the generalized matching pennies problem. Thus, we do not obtain an improved advice complexity lower bound for the MTS problem on a worst-case metric space. We remark that the metric space and the type of tasks used to obtain the $\Omega(\log N)$ lower bound for algorithms with $o(n)$ bits of advice in [48] is of such a nature that the lower bound does not imply a similar lower bound for e.g. k-server or paging.

As with the k-server problem, combining Theorem 1 with the results of Bartal et al. [15] yields a lower bound of $\Omega(\log N/\log^2 \log N)$ for algorithms with sublinear advice on arbitrary metric spaces. No advice complexity lower bounds for the MTS problem on arbitrary metric spaces were previously known.

Contrary to the k-server problem, almost tight lower and upper bounds on the competitive ratio of randomized algorithms without advice are known. In particular, there exists a randomized $O(\log^2 N \log \log N)$-competitive MTS-algorithm for arbitrary metric spaces [50]. However, there is still a small gap of $\Theta(\log N \log \log N)$ between the best known upper and lower bound for randomized algorithms without advice. Theorem 2 shows that finding the best possible competitive ratio of deterministic algorithms with advice complexity $o(n)$ would close this gap for finite metric spaces. We remark that it is possible to achieve a competitive ratio of $O(\log N)$ using n bits of advice [48].

6.4 List update

The list update problem on a finite list can be modeled as a finite LTS and is therefore compact and Σ-repeatable. In [30], a lower bound of $15/14$ for list update algorithms with sublinear advice was shown by a reduction from binary string guessing. We improve this lower bound to $3/2$ by combining the fact that list update is Σ-repeatable with a lower bound of $3/2$ for randomized algorithms without advice due to Teia [91].

The best known randomized list update algorithm is the 1.6-competitive COMB algorithm [4]. By Theorem 2, achieving a competitive ratio of $1.5 \leq c < 1.6$ with $o(n)$ bits of advice would improve on this upper bound.

6.5 Bipartite matching

We consider the (unweighted) bipartite matching problem [69]. The initial state of the problem is a number $l \in \mathbb{N}$. A sequence of requests is defined in terms of a bipartite graph $G = (L, R, E)$ (where L and R are the parts of the bipartition of the vertex set and E is the set of edges) satisfying $|L| = l$. The vertices of L are called the *offline* vertices. A request is a vertex $v \in R$ together with all edges between L and r. As soon as a vertex arrives, it must irrevocably be matched to an incident (and unmatched) vertex in L, if possible. The goal is to match as many vertices as possible. Note that the number of requests, n, equals the size of R. It is usually assumed that $n = |R| = \Theta(l)$.

Note that the size l of L is part of the initial state of the problem. In particular, the additive constant in the definition of the competitive ratio cannot depend on l since l is part of the input. This is necessary to avoid trivializing the problem, since $\texttt{OPT}(\sigma) \leq l$ for all inputs σ.

Let P be the bipartite matching problem. We will show that P is Σ-repeatable. Given an input $\sigma^* = (\sigma_1; \ldots; \sigma_r)$ for P^*_Σ where all rounds have l offline vertices, we define the input $g(\sigma^*)$ for P as follows: The input $g(\sigma^*)$ will have $r \cdot l$ offline vertices[8]. We partition these vertices into r groups such that the ith group contains l vertices. The requests of σ_1 will be revealed using the first group of offline vertices, then the requests of σ_2 will be revealed using the second group of offline vertices, and so on. It is easy to see that this reduction shows that bipartite matching is Σ-repeatable with parameters $(0, 0, 0)$.

Previously, it was known that $\lceil \log(n!) \rceil = \Omega(n \log n)$ bits of advice were necessary to be strictly 1-competitive [83]. A classical result of Karp et al. says that there exists a randomized $e/(e-1)$-competitive algorithm, and that no randomized algorithm (without advice) can achieve a better competitive ratio [69]. We get from this and Theorem 1 that no randomized algorithm with advice complexity $o(n)$ can achieve a competitive ratio better than $e/(e-1)$, and that this is tight. Shortly after our work, it was shown by Dürr et al. (using a streaming algorithm from [46]) that for every constant $\varepsilon > 0$, there exists a $(1+\varepsilon)$-competitive algorithm using $O(n)$ bits of advice [45].

6.6 Reordering buffer management

It is easy to show that the reordering buffer management problem with a buffer of size k is Σ-repeatable (see Lemma 16). In [1], Adamaszek el al. give a lower bound of $\Omega(\log \log k)$ for randomized online algorithms without advice. Since reordering buffer management is Σ-repeatable (and since the lower bound in [1] satisfies the conditions of Theorem 1), we get the same lower bound for algorithms with advice complexity $o(n)$. In [8], a randomized $O(\log \log k)$-competitive algorithm is obtained. Using derandomization, we see that there exists a deterministic $O(\log \log k)$-competitive algorithm with $O(\log n)$ bits of advice. Finally, it is shown in [2] that for every $\varepsilon > 0$, there is a deterministic $(1+\varepsilon)$-competitive algorithm reading $O(n)$ bits of advice and that $\Omega(n \log k)$ bits of advice is necessary (and obviously sufficient) to be 1-competitive. Thus, the asymptotic number of advice bits needed to achieve a given asymptotic competitive ratio is now fully determined for randomized algorithms.

6.7 Dynamic binary search trees

We will consider the dynamic binary search tree problem from the perspective of competitive analysis. Several models exists, most of which are equivalent up to constant factors (see e.g. [92]). The model that we use is as follows: The goal is to create a binary search tree (BST) which supports searches for keys from a static set $[N]$. An input is a sequence of keys $\sigma = (x_1, \ldots, x_n)$ such that $x_i \in [N]$ for each $1 \leq i \leq n$ (so that searches are always successful). A BST-algorithm maintains a pointer. The pointer can be moved to its left child, its right child, or its parent, all at unit cost. Also, the BST-algorithm can perform a rotation on the current node and its parent at unit cost. When x_i is revealed, the BST-algorithm must make a sequence of unit-cost operations such that the node containing the key x_i is at some point visited. We assume that the pointer starts in the node where it ended after the search for the previous key x_{i-1} in the input sequence. Finally, we assume that initializing the pointer (at the last node visited) is a unit-cost operation which must

[8] Since the number of offline vertices is part of the input, the reduction g from P^*_Σ to P is allowed to change it.

always be performed at the beginning of a search. This assumption ensures that the cost of finding a key is at least the number of nodes visited.

Demaine et al. has shown that there exist $O(\log \log N)$-competitive binary search trees [42]. The well-known dynamic optimality conjecture claims that Sleater and Tarjan's splay trees [89] are in fact $O(1)$-competitive. We show in Lemma 15 in Appendix D that the dynamic binary search tree problem (as it is defined above) can be modeled as a lazy task system. This implies that if there exists a BST with advice complexity $o(n)$ where n is the number of searches (and not the number of keys), then there also exists a randomized $O(1)$-competitive BST without advice. We remark that it has previously been shown to be possible to (effectively) convert a BST with advice complexity $O(1)$ into a deterministic BST with essentially the same performance guarantee [42].

6.8 Sleep states management

The 2-sleep states management problem (2-SSM, see [64, 21]) can, obviously, be modeled as a generalized task system, and so it is Σ-repeatable. As observed in [64], the 2-SSM problem is essentially a repeated ski-rental problem. Thus, it follows from [67] that no randomized algorithm without advice can be better than $e/(e-1)$-competitive. Furthermore, this lower bound satisfies the conditions of Theorem 1. Thus, we get a lower bound of $e/(e-1)$ for algorithms with advice complexity $o(n)$ for 2-SSM. The previous best lower bound for algorithms with sublinear advice was 7/6 and was obtained in [21]. We note that [67] also gives a randomized algorithm with a competitive ratio of $e/(e-1)$. Furthermore, in [21], it is shown that it is possible for a deterministic algorithm to achieve a competitive ratio arbitrarily close to $e/(e-1)$ using $O(1)$ bits of advice. Since n bits of advice trivially suffice to be optimal for 2-SSM, the asymptotic advice complexity of 2-SSM is now fully determined.

6.9 One-dimensional unit clustering

We consider the one-dimensional online unit clustering problem (see [34, 49]). We claim that this problem is Σ-repeatable with parameters $(0, 0, 0)$. Let $\sigma^* = (\sigma_1; \ldots; \sigma_r)$ be an input for the repeated version of the problem. For $1 \leq i \leq r$, let $x^i_{\min}$ be the left-most point of σ_i, and let $x^i_{\max}$ be the right-most point of σ_i. Thus, for every point $x \in \sigma_i$, it holds that $x \in [x^i_{\min}; x^i_{\max}]$. Let $L_i = \left|x^i_{\max} - x^i_{\min}\right|$. Furthermore, let $a_1 = 0, b_1 = L_1$ and define recursively $a_i = b_{i-1} + 2$ and $b_i = a_i + L_i$ for $2 \leq i \leq r$. We construct an input σ for the online unit clustering problem as follows: The points of σ_i are shifted so that they all fall within the interval $[a_i, b_i]$. This is possible since $b_i - a_i = L_i$. Furthermore, for $i \neq i'$, a cluster cannot cover points from both $[a_i, b_i]$ and $[a_{i'}; b_{i'}]$. It follows easily that this reduction satisfies the conditions of Definition 6.

It is shown in [49] that every randomized algorithm is at least 3/2-competitive. This is done by showing that for every $\varepsilon > 0$, there exists a hard input distribution with finite support such that no deterministic algorithm can be better than $(3/2 - \varepsilon)$-competitive against this distribution. Note that the number of possible requests for the unit clustering problem is uncountable. However, the lower bound of 3/2 is proved using input distributions with finite support. Thus, combining the lower bound of [49] with Theorem 1, it follows that an algorithm with $o(n)$ bits of advice must be at least 3/2-competitive. We remark that the best known online algorithm for one-dimensional unit clustering is a 5/3-competitive deterministic algorithm [47].

6.10 Max-SAT

The online Max-SAT problem (see [10]) is Σ-repeatable with parameters $(0, 0, 0)$. Indeed, given an input $\sigma^* = (\sigma_1; \ldots; \sigma_r)$ for the repeated version of the problem, we can construct an input σ of Max-SAT by using a fresh set of variables for each round σ_i (that is, for $i \neq i'$, the set of variables used in round i is disjoint from the set of variables used in round i'). It is easy to see that this reduction satisfies the conditions of Definition 6.

It is shown in [10] that no randomized algorithm can achieve a competitive ratio better than $3/2$. This is done using a family of hard input distributions which are compatible with Theorem 1. Thus, we get that an algorithm with advice complexity $o(n)$ must have a competitive ratio of at least $3/2$. We remark that a $3/2$-competitive randomized algorithm (without advice) is given in [10].

6.11 Metric matching

In the metric matching problem (see [65, 70]), we are given as initial state a metric space (M, d) and k servers $S \subseteq M$ with a pre-specified location[9]. A request r_i is a point $r_i \in M$ in the metric space. When a request r_i arrives, an online algorithm must irrevocably match the request to a server $s \in S$ which has not previously been matched to another request. The algorithm incurs a cost of $d(r_i, s)$ for matching the request r_i to the server s. Note in particular that the number of requests is (at most) k.

We will show that the metric matching problem is Σ-repeatable with parameters $(0, 0, 0)$. To this end, we will use that the underlying metric space and the pre-specified servers are parts of the input (and not fixed as parameters of the problem). As with bipartite matching, the problem would trivialize if the additive constant was allowed to depend on k and the metric space (since the cost of `OPT` is never more that $k\Delta$ where Δ is the diameter). Note that this is very different from the k-server problem.

Given an input $\sigma^* = (\sigma_1; \ldots; \sigma_r)$ for the repeated version of the problem, we can construct an input σ of metric matching as follows: Denote by (M', d') and $S' \subseteq M'$ the metric space and set of servers used as initial state for σ^* (recall that all of $\sigma_1, \ldots, \sigma_r$ by definition are required to have the same initial state). We define a new metric space (M, d) which contains r subspaces, $(M_1, d_1), \ldots, (M_r, d_r)$, each of which is a copy of (M', d'), in such a way that the distance between two points in two different subspaces is at least B for some sufficiently large $B \in \mathbb{R}$. The requests of σ are simply the requests of $\sigma_1, \ldots, \sigma_r$ such that the requests of σ_i are points in (M_i, d_i). By choosing B to be sufficiently large, we can ensure that if an algorithm ever matches a request in σ_i to a server in a subspace other than (M_i, d_i), then that algorithm incurs a cost larger than any algorithm which only matches requests to servers in their corresponding subspace. With such a choice of B, it is easy to see that the conditions of Definition 6 are satisfied.

For general metric spaces, there exists an $O(\log^2 k)$-competitive randomized algorithm [12]. Also, it is known that any randomized algorithm must be $\Omega(\log k)$-competitive [81]. Note that $\Omega(\log k)$ is not constant with respect to the input length. However, together with Theorem 1, it does imply that no algorithm with sublinear advice can be $O(1)$-competitive.

From the discussion at the end of Section 5.1, it is easy to see that the metric matching problem is compact (on finite metric spaces). Indeed, in this problem, there are k servers placed in a metric space. Each server can be matched to at most one request (and vice versa). But this means that

[9] Note that S is a multisubset of M since more than one server could be placed at the same point in M.

for every fixed metric space and fixed set of k servers, the length of the input never exceeds k. This allow (reusing the notation from the discussion) us to construct a single algorithm $\mathtt{R}$ which on inputs with k servers run the appropriate algorithm $\mathtt{R}_k$.

An interesting special case is the line metric. An $O(\log k)$-competitive randomized algorithm is known [58], but even for deterministic algorithms, the best lower bound is only 9.001 [56]. Therefore, it has been conjectured that there exists an $O(1)$-competitive randomized algorithm (see e.g. [66]). If instead of the real line we consider the problem on finite metric discretizations of the line, then the metric matching problem becomes compact. Thus, it follows from Theorem 2 that this conjecture is equivalent to the conjecture that there exists a deterministic $O(1)$-competitive algorithm with sublinear advice for the line metric.

6.12 Examples of non-repeatable online problems

Some examples of classical online problems which are provably not Σ-repeatable includes bin packing (see however Section 8.4.1) and makespan minimization (on identical machines). For both of these problems, it is has been shown that sublinear advice allows an algorithm to achieve a competitive ratio which is provably better than the competitive ratio of any randomized online algorithm without advice [31, 3]. Since the lower bounds for algorithms without advice satisfy the conditions of Theorem 1, this proves that the problems are not Σ-repeatable.

7 Constructive version of Theorem 2 via online learning

As mentioned, the proof of Theorem 2 is non-constructive. It does not show how to obtain the randomized algorithm given an algorithm with sublinear advice. In this section, we show how to easily obtain a constructive version of an important special case of Theorem 2 by combining a variant of the Weighted Majority Algorithm from online learning with the Reset Lemma (Lemma 9). Recall that Theorem 2 holds for all compact and Σ-repeatable problems. However, the most interesting class of problems known to be compact are those that can be formulated as a (finite) LTS. We will show how to obtain an algorithmic version of Theorem 2 for lazy task systems.

Theorem 7. *Let* P *be a minimization problem which can be modeled as a finite lazy task system, let* $c \geq 1$ *be a constant and let* $\mathtt{ALG}$ *be a deterministic* c*-competitive* P*-algorithm with advice complexity* $o(n)$*. Then, for every* $\varepsilon > 0$*, there exists a randomized* $(c+\varepsilon)$*-competitive algorithm without advice.*

Proof. By assumption, P can be modeled as a finite LTS $(\mathcal{S}, \mathcal{T})$. Since $(\mathcal{S}, \mathcal{T})$ is finite, it has a finite max-cost $\Delta \in \mathbb{R}$. Let $\varepsilon_1 > 0$ be arbitrary and let b be the advice complexity of the c-competitive P-algorithm $\mathtt{ALG}$. We want to apply the Reset Lemma. To this end, let $n' \in \mathbb{N}$. Note that for inputs of length at most n', the algorithm $\mathtt{ALG}$ can be converted into $m = 2^{b(n')}$ deterministic algorithms, $\mathtt{ALG}_1, \ldots, \mathtt{ALG}_m$, without advice. Following Blum and Burch [20] (see also Section 1.4), we use the algorithm Hedge [55] to combine the m algorithms into a single randomized algorithm, $\mathtt{R}_{n'}$, such

that for every input σ of length at most n', it holds that (cf. (2)):

$$\begin{aligned}
\mathbb{E}[\mathtt{R}_{n'}(\sigma)] &= (1+2\varepsilon_1)\cdot \min_{1\le i\le m} \mathtt{ALG}_i(\sigma) + \left(\frac{7}{6}+\frac{1}{\varepsilon_1}\right)\Delta \ln m \\
&= (1+2\varepsilon_1)\cdot \mathtt{ALG}(\sigma) + \left(\frac{7}{6}+\frac{1}{\varepsilon_1}\right)\Delta\, b(n')\ln(2). \\
&\le (1+2\varepsilon_1)\cdot (c\cdot \mathtt{OPT}(\sigma)+\alpha_0) + \left(\frac{7}{6}+\frac{1}{\varepsilon_1}\right)\Delta\, b(n')\ln(2) \\
&= (1+2\varepsilon_1)\cdot (c\cdot \mathtt{OPT}(\sigma)) + (1+2\varepsilon_1)\alpha_0 + \left(\frac{7}{6}+\frac{1}{\varepsilon_1}\right)\Delta\, b(n')\ln(2).
\end{aligned}$$

Define $\alpha : \mathbb{N} \to \mathbb{R}_{\ge 0}$ as follows:

$$\alpha(n) = (1+2\varepsilon_1)\alpha_0 + \left(\frac{7}{6}+\frac{1}{\varepsilon_1}\right)\Delta\, b(n)\ln(2). \tag{40}$$

By assumption, $b(n) \in o(n)$. Furthermore, ε_1, α_0, and Δ are constants (with respect to n). Thus, $\alpha(n) \in o(n)$.

We now have that for each $n' \in \mathbb{N}$, there exists a randomized algorithm $\mathtt{R}_{n'}$ such that $\mathbb{E}[\mathtt{R}_{n'}(\sigma)] \le (1+2\varepsilon_1)c\cdot\mathtt{OPT}(\sigma)+\alpha(n')$ for every input σ of length $|\sigma| \le n'$. Thus, since $\alpha(n) \in o(n)$, the Reset Lemma (Lemma 9) shows that for every $\varepsilon_2 > 0$, there exists a randomized algorithm $\mathtt{R}$ without advice such that $\mathtt{R}$ is $((1+2\varepsilon_1)c+\varepsilon_2)$-competitive on inputs of arbitrary length. Recall that c is a constant. Choosing $\varepsilon_1 = \varepsilon/(4c)$ and $\varepsilon_2 = \varepsilon/2$ yields a randomized $(c+\varepsilon)$-competitive algorithm without advice.

□

8 Repeated matrix games

One of the most successful techniques for proving advice complexity lower bounds has been via reductions from the string guessing problem [48, 24]. In this section, we show that string guessing is a special case of a more general online problem where the algorithm and the adversary repeatedly play a zero-sum game. Applications to bin packing and paging are given.

We first fix the notation. Let $q \ge 2$ and let $[q] = \{1, 2, \ldots, q\}$. Let $A \in \mathbb{R}_{\ge 0}^{q\times q}$ be a quadratic $q \times q$ matrix with non-negative real entries. Such a matrix defines a finite two-player zero-sum game: Player I, the "row-player", has q pure strategies, one for each row of A. Player II, the "column-player", has q pure strategies, one for each column of A. Simultaneously, Player I chooses a row $x \in [q]$ and Player II chooses a column $y \in [q]$. The payoff function for player I is $A(x,y)$ and the payoff function for player II is $-A(x,y)$. If player I uses the mixed strategy[10] $\mu : [q] \to [0,1]$ and player II uses the mixed strategy $\nu : [q] \to [0,1]$ then the expected payoff of player I is defined as

$$A(\mu,\nu) = \mathop{\mathbb{E}}_{\substack{x\sim\mu\\ y\sim\nu}}[A(x,y)] = \sum_{x=1}^{q}\sum_{y=1}^{q} A(x,y)\mu(x)\nu(y).$$

Since the game defined by A is a finite two-player zero-sum game, it follows from von Neumann's minimax theorem that it has a *value* $V \ge 0$ such that

$$V = \max_{\mu}\min_{\nu} A(\mu,\nu) = \min_{\nu}\max_{\mu} A(\mu,\nu). \tag{41}$$

[10] A mixed strategy is a probability distribution over rows (or columns).

Here, the max is over all mixed strategies μ for player I, and the min is over all mixed strategies ν for player II. The value V (along with optimal strategies) is efficiently computable using linear programming. As a corollary to (41), sometimes known as Loomis' lemma, it follows that there exists an optimal mixed strategy, $\mu^\star$, for player I and an optimal mixed strategy, $\nu^\star$, for player II such that $V = \max_{x\in[q]} A(x,\nu^\star) = \min_{y\in[q]} F(\mu^\star, y)$.

The game defined above is sometimes called a *one-shot* game. Given such a one-shot game, one may define a *repeated game* where the two players repeatedly play the one-shot game. We will now define an online version of this repeated game.

Definition 12. Let $q \in \mathbb{N}$. A *repeated matrix game (RMG)* is an online problem defined by a cost matrix $A \in \mathbb{R}_{\geq 0}^{q\times q}$. The input $\sigma = (n, x_1, \ldots, x_n)$ consists of an integer n and a sequence of n characters from $[q]$. Round $1 \leq i \leq n$ of the game proceeds according to the following rules:

1. If $i = 1$, the algorithm learns n. If $i > 1$, the algorithm learns the input character x_{i-1}.
2. The algorithm answers $y_i = f_i(n, x_1, \ldots, x_{i-1}) \in [q]$, where f_i is a function defined by the algorithm.

The cost of the output $\gamma = y_1 \ldots y_n$ computed by the algorithm is $\sum_{i=1}^{n} A(x_i, y_i)$.

Note that in a repeated matrix game, the adversary is the row-player and the algorithm is the column-player. We will not consider the competitive ratio of a RMG-algorithm but instead the total cost incurred by the algorithm. Also, following previous work [48, 24, 29] on string guessing problems, we slightly abuse notation by using n to denote the number of input characters and not the number of requests, which is $n+1$ (in the standard model of online problems, it is not possible to avoid this extra round since the ith input character must be revealed strictly later than when the algorithm has to output its ith character). Also, note that RMG does not have any initial state (n is not the initial state but a request that must be answered by the algorithm).

For a RMG with cost matrix A, there exists a deterministic algorithm which incurs a cost of at most $\min_{y\in[q]} \max_{x\in[q]} A(x,y)$ in each round by playing according to an optimal pure strategy for the column player. Clearly, no deterministic algorithm (without advice) can do better than this. On the other hand, a randomized algorithm can ensure an expected cost of at most V in each round by playing according to an optimal mixed strategy for the column player. A small amount of advice suffices to derandomize this strategy.

Theorem 8. *Let V be the value of the (one-shot) two-player zero-sum game defined by the matrix $A \in \mathbb{R}_{\geq 0}^{q\times q}$. Then, there exists a randomized algorithm $\texttt{R}$ without advice for the RMG with cost matrix A such that $\mathbb{E}[\texttt{R}(\sigma)] \leq Vn$ for every input σ. Furthermore, there exists a deterministic algorithm, $\texttt{ALG}$, reading $\lceil \log q \rceil$ bits of advice such that $\texttt{ALG}(\sigma) \leq Vn$ for every input σ.*

Proof. Consider the one-shot game defined by A. From the minimax theorem, we know that there exists an optimal mixed strategy $\nu^* : [q] \to [0,1]$ for the column-player such that for each $x \in [q]$, it holds that $A(x,\nu^*) = \mathbb{E}_{y\sim\nu^*}[A(x,y)] = \sum_{y=1}^{q} A(x,y)\nu^*(y) \leq V$. We will now define the algorithm $\texttt{R}$ for the RMG with cost matrix A. Before the first request arrives, the algorithm $\texttt{R}$ selects a character, $Y \in [q]$, at random according to ν^* (note that Y is a random variable). In each round, $\texttt{R}$ answers Y. Let $1 \leq i \leq n$ and let $x_i \in [q]$ be the input character in round i. We know that the expected cost $A(x_i, \nu^*)$ incurred by $\texttt{R}$ in round i is at most V. Thus, by linearity of expectation, $E[\texttt{R}(\sigma)] \leq Vn$ for every input σ.

We can convert R into a deterministic algorithm ALG with advice. The oracle looks at the input σ and determines the best choice, y, of the random variable Y for this particular input (that is, the oracle finds the character y minimizing $\sum_{i=1}^{n} A(x_i, y)$). It writes y onto the advice tape using $\lceil \log q \rceil$ bits of advice. The algorithm ALG learns y from the advice and answers y in each round. It follows that $\texttt{ALG}(\sigma) \leq \mathbb{E}[\texttt{R}(\sigma)] \leq Vn$ for every input σ. □

8.1 Proof of Theorem 3

We will show that linear advice is required to ensure a cost of at most $(V-\varepsilon)n$ if $\varepsilon > 0$ is constant. This result follows from Lemma 2 and the fact that all RMGs are Σ-repeatable. However, it is much more natural to prove this result directly using our direct product theorem (Theorem 7), since the result is in fact an easier version of Lemma 2 where each round consists of exactly one request. Also, there is a slight technical problem with using the fact that a RMG is Σ-repeatable if one wants a non-asymptotic lower bound: A RMG-input must consist of at least two requests (since the number of requests is $n+1$, where n is the number of input characters). By our definitions of Σ-repeatable online problems and our statement of Lemma 2, this will result in a lower bound which is unnecessarily a factor of two lower than it should be. Therefore, we give a direct proof of Theorem 3 instead of using Lemma 2. For convenience, we restate Theorem 3 before giving the proof.

Theorem 3. *Let ALG be an algorithm for the RMG with cost matrix A. Furthermore, let V be the value of the (one-shot) two-person zero-sum game defined by A and let $0 < \varepsilon \leq V$ be a constant. If $\mathbb{E}[\texttt{ALG}(\sigma)] \leq (V-\varepsilon)n$ for every input σ of length n, then ALG must read at least*

$$b \geq \frac{\varepsilon^2}{2\ln(2) \cdot \|A\|_\infty^2} n = \Omega(n) \tag{1}$$

bits of advice.

Proof of Theorem 3. Since the value of the game defined by A is V, we know that there exists a probability distribution μ^* over characters from $[q]$ such that for every $y \in [q]$, it holds that $A(\mu^*, y) \geq V$.

Let $n \in \mathbb{N}$. Define $p^n \colon \{(n, x_1, \ldots, x_n) | x_i \in [q]\} \to [0,1]$ to be the RMG-input distribution which maps $(n, x_1, \ldots, x_n)$ into $\mu^*(x_1)\mu^*(x_2)\cdots\mu^*(x_n)$. Thus, in each of the n rounds, we draw independently an input character from $[q]$ according to μ^*. The ith round cost function simply assigns a cost of $A(x_i, y_i)$ to the algorithm where x_i is the input character and y_i is the answer produced by the algorithm in round i. Together with these cost-functions, the input distribution p^n is an n-round input distribution. Note that the ith round input distribution p_i is simply $p_i = \mu^*$.

Let ALG be a deterministic RMG-algorithm reading at most b bits of advice on inputs of length at most n. Fix $1 \leq i \leq n$ and $w \in \mathcal{W}_i$. Let $d = D_{\mathrm{KL}}(p_{i|w} \| p_i)$. In order to apply Theorem 1, we need a lower bound on $\mathbb{E}[\mathrm{cost}_i(\texttt{ALG}) | W_i = w]$ in terms of d. By Pinsker's inequality (9),

$$\|p_{i|w} - p_i\|_1 \leq \sqrt{d \cdot \ln 4}. \tag{42}$$

Let $h(d) = \sqrt{d \cdot \ln 4}$. For every $w \in \mathcal{W}_i$ where $\Pr[W_i = w] > 0$, the algorithm ALG outputs a fixed character $y_w \in [q]$ in round i and incurs an expected cost of $\mathbb{E}_{x \sim p_{i|w}}[A(x, y_w)]$. Obviously, the cost

incurred by ALG in any round is never larger than $\|A\|_\infty := \max_{x,y} A(x,y)$. Thus,

$$\begin{aligned}
\mathbb{E}[\mathrm{cost}_i(\texttt{ALG})|W_i = w] &= \mathbb{E}_{x\sim p_{i|w}}[A(x,y_w)] \\
&\geq \mathbb{E}_{x\sim p_i}[A(x,y_w)] - \|A\|_\infty \cdot h(d) && \{\text{By (42) and (8)}\} \\
&= \mathbb{E}_{x\sim \mu^*}[A(x,y_w)] - \|A\|_\infty \cdot h(d) \\
&\geq V - \|A\|_\infty \cdot h(d).
\end{aligned}$$

Define $f(d) = V - \|A\|_\infty \cdot h(d)$. Since f is convex and decreasing, Theorem 7 yields

$$\mathbb{E}_{\sigma\sim p^n}[\texttt{ALG}(\sigma)] = \mathbb{E}[\mathrm{cost}(\texttt{ALG})] \geq n f(b/n) = n\left(V - \|A\|_\infty \sqrt{\frac{b\cdot \ln 4}{n}}\right). \tag{43}$$

It follows from Yao's principle that if a randomized algorithm on every input of length n reads at most b bits of advice and incurs an expected cost of at most $(V-\varepsilon)n$, then

$$(V-\varepsilon)n \geq n\left(V - \|A\|_\infty \sqrt{\frac{b\cdot \ln 4}{n}}\right).$$

A straightforward calculation completes the proof of the theorem.

□

Since $\varepsilon \leq V \leq \|A\|_\infty$, we have $\frac{\varepsilon^2}{2\ln(2)\|A\|_\infty^2} \leq \frac{1}{2\ln(2)} < 0.722$. Thus, the lower bound (1) is never larger than $0.722n$ and will often be much lower. In what follows, we will obtain better bounds than (1) for some specific cost matrices.

8.2 Guessing strings and matching pennies

For alphabets of size $q = 2$, string guessing with known history [48, 24] is the repeated matrix game with cost matrix

$$A_2 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

More generally, the q-SGKH problem [48, 24] is the repeated matrix game with cost matrix $A_q = J_q - I_q$ where J_q is the all-ones matrix and I_q the identity matrix, both of size $q \times q$. That is,

$$A_q = \begin{pmatrix} 0 & 1 & 1 & \dots & 1 \\ 1 & 0 & 1 & \dots & 1 \\ 1 & 1 & 0 & \dots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & 0 \end{pmatrix}.$$

Emek et al. introduced the generalized matching pennies problem [48] which is equivalent to the string guessing problem, except that a dummy cost of $1/\tau$ is added to every entry of A_q for some integer τ. Using the high-entropy technique, they obtained lower bounds valid for alphabets of size $q \geq 4$ and certain ranges of the advice read and the cost incurred. Böckenhauer et al. later obtained tight (up to an additive lower order term of order $O(\log n)$) lower and upper bounds on the deterministic advice complexity of the string guessing problem [24]. Since their introduction, a

large number of advice complexity lower bounds have been obtained by reductions from the string guessing or generalized matching pennies problem. This includes lower bounds for bin packing [31, 6], list update [30], metrical task systems [48], reordering buffer management [2], set cover [73], and several other online problems.

It is easy to see that the value V_q of the game defined by the matrix A_q is $V_q = (q-1)/q$, and that the optimal mixed strategy for both players is to select a character uniformly at random. Thus, it follows immediately from Theorem 3 that $\Omega(n)$ bits of advice are needed for an algorithm to ensure a cost of at most $(\frac{q-1}{q} - \varepsilon)n$. Note how Theorem 3 makes it a trivial task to obtain this result. For several applications, the lower bound of $\Omega(n)$ is all that one needs. However, if one seeks a better lower bound for larger alphabets, or if the exact coefficient of the higher-order term is important, a better lower bound is needed. In this section, we will show how to reprove the tight lower bounds due to Böckenhauer et al. [24] using our direct product theorem. We believe that our technique yields a very intuitive proof of the lower bound. In particular, it explains nicely why the q-ary entropy function appears in the tight lower bound. Also, we slightly strengthen the lower bound of [24] by showing that it also applies to randomized algorithms with advice (this also follows from our derandomization results in Appendix A).

Theorem 9. *Let $q \geq 2$ and let $0 < \alpha < (q-1)/q$. A randomized q-SGKH algorithm which on every input of length n has an expected cost of at most αn must read at least*

$$b \geq K_{1/q}(1-\alpha)n = (1 - h_q(\alpha))n \log_2 q$$

bits of advice.

Proof. Let $n \in \mathbb{N}$. For $1 \leq i \leq n$, the input character x_i is chosen uniformly at random from $[q]$. This gives rise to an n-round input distribution where each round consists of exactly one request. The ith round cost-function cost_i simply assigns a cost of 0 to an algorithm in round i if the algorithm guesses correctly and a cost of 1 otherwise.

Fix a deterministic algorithm ALG which reads at most b bits of advice on inputs of length n. Let $1 \leq i \leq n$ and let $w \in \mathcal{W}_i$ be any possible history in round i. Let $d = D_{KL}(p_{i|w} \| p_i)$. Then, by Lemma 1, for every $x \in [q]$ it holds that $p_{i|w}(x) \leq K^{-1}_{1/q,r}(d)$. Thus,

$$\mathbb{E}[\text{cost}_i(\texttt{ALG}) | W_i = w] \geq 1 - K^{-1}_{1/q,r}(d),$$

since no matter which character that ALG chooses in round i, the probability of this character being wrong (and ALG therefore incurring a cost of 1) is at least $1 - K^{-1}_{1/q,r}(d)$. Let $f(d) = 1 - K^{-1}_{1/q,r}(d)$. Note that f is convex and decreasing. By Theorem 7, we therefore get that $\mathbb{E}[\text{cost}(\texttt{ALG})] \geq nf(n^{-1}b)$. We will now calculate how large b needs to be in order to ensure that the expected cost is at most αn.

$$\begin{aligned}
&\alpha n \geq nf(b/n) \\
&\Rightarrow \alpha \geq 1 - K^{-1}_{1/q,r}(b/n) \\
&\Rightarrow K^{-1}_{1/q,r}(b/n) \geq 1 - \alpha \\
&\Rightarrow b/n \geq K_{1/q}(1-\alpha) && \{K_{1/q} \text{ is increasing on } [1/q, 1]\} \\
&\Rightarrow b \geq nK_{1/q}(1-\alpha) = n(1 - h_q(\alpha)) \log(q).
\end{aligned}$$

Thus, any deterministic algorithm needs to read at least $b \geq K_{1/q}(1-\alpha)n$ bits of advice in order to ensure an expected cost of at most αn against this n-round input distribution. By Yao's principle, we get that any randomized algorithm must read at least $K_{1/q}(1-\alpha)n$ bits of advice in order to guarantee that it always incurs an expected cost of at most αn. □

As mentioned, it is shown in [24] that Theorem 9 is tight up to (additive) lower order terms. Note that Theorem 9 shows that we need $\Omega(n \log q)$ bits to ensure a cost of at most $(\frac{q-1}{q} - \varepsilon)n$, whereas Theorem 3 only gives a lower bound of $\Omega(n)$.

8.3 Anti-string guessing

The *anti-string guessing problem (Anti-q-SG)* over an alphabet of size q is the repeated matrix game with cost-matrix I_q (the identity matrix of size $q \times q$). Thus, the goal of the algorithm is to output a character *different* from the input character. The value of the one-shot game defined by I_q is $1/q$. For both players, the optimal strategy is to select a character uniformly at random. Theorem 3 therefore implies that an online algorithm needs $\Omega(n)$ bits of advice in order to achieve a cost of at most $(1/q-\varepsilon)n$ on every input of length n. As was the case for the original string guessing problem, Theorem 3 makes it trivial to obtain this result. Note that for $q = 2$, string guessing and anti-string guessing are equivalent. However, one interesting aspect of anti-string guessing is that it becomes easier as q grows.

In this section, we will show that a direct application of our information theoretical direct product theorem yields lower bounds for anti-string guessing which are not just asymptotically tight, but are tight up to an additive $O(\log n)$ term.

Theorem 10. *Let $q \geq 2$ and let $0 < \alpha < 1/q$. A randomized Anti-q-SG algorithm which on every input of length n incurs an expected cost of at most αn must read at least*

$$b \geq K_{1/q}(\alpha)n = (1 - h_q(1-\alpha))n \log_2 q$$

bits of advice.

Proof. Let $n \in \mathbb{N}$. For $1 \leq i \leq n$, the input character x_i is chosen uniformly at random from $[q]$. This gives rise to an n-round input distribution where each round consists of exactly one request. The ith round cost-function cost_i simply assigns a cost of 0 to an algorithm in round i if the algorithm outputs something else than the input character and a cost of 1 if the algorithm outputs the input character.

Fix a deterministic algorithm `ALG` which reads at most b bits of advice on inputs of length n. Let $1 \leq i \leq n$ and let $w \in \mathcal{W}_i$ be any possible history in round i. Let $d = D_{KL}(p_{i|w} \| p_i)$. Then, by Lemma 1, the smallest probability assigned to a single character by $p_{i|w}$ is at least $K^{-1}_{1/q,l}(d)$. Thus, $\mathbb{E}[\text{cost}_i(\texttt{ALG}) | W_i = w] \geq K^{-1}_{1/q,l}(d)$, since no matter which character that `ALG` chooses in round i, the probability of this character being equal to the input character is at least $K^{-1}_{1/q,l}(d)$. Let $f(d) = K^{-1}_{1/q,l}(d)$. Note that f is convex and decreasing. By Theorem 7, we have that $\mathbb{E}[\text{cost}(\texttt{ALG})] \geq nf(n^{-1}b)$. We will now calculate how large b needs to be in order to ensure

that the expected cost is at most αn:

$$\begin{aligned}
&\alpha n \geq n f(b/n) \\
&\Rightarrow \alpha \geq K^{-1}_{1/q,l}(b/n) \\
&\Rightarrow K_{1/q}(\alpha) \leq b/n && \{K_{1/q} \text{ is decreasing on } [0, 1/q]\} \\
&\Rightarrow b \geq n K_{1/q}(\alpha) = n(1 - h_q(1-\alpha)) \log(q).
\end{aligned}$$

Thus, any deterministic algorithm needs to read at least $b \geq K_{1/q}(\alpha)n$ bits of advice in order to ensure an expected cost of at most αn against this n-round input distribution. By Yao's principle, we get that any randomized algorithm must read at least $K_{1/q}(\alpha)n$ bits of advice in order to guarantee that its expected cost is never more than αn.

□

Upper bound for anti-string guessing. It follows from Theorem 8 that it is possible to ensure an expected cost of at most n/q for Anti-q-SG by simply answering randomly in each round (or by using $\lceil \log q \rceil$ bits of advice to communicate the character appearing least frequently in the input). We will now show how a simple probabilistic argument yields an upper bound on the advice needed to ensure a cost less than n/q. This upper bound shows that Theorem 10 is essentially tight.

Theorem 11. *Let $q \geq 2$ and let $0 < \alpha < 1/q$. There exists a deterministic algorithm for Anti-q-SG which on every input of length n incurs a cost of at most αn and which reads*

$$b = K_{1/q}(\alpha)n + O(\log n + \log\log q) = (1 - h_q(1-\alpha))n \log_2 q + O(\log n + \log\log q)$$

bits of advice.

Proof. See Appendix E.2. □

8.3.1 Application to paging

We have already shown using Theorem 1 that $\Omega(n)$ bits of advice are needed to be better than H_k-competitive for paging. However, as usual, lower bounds obtained from Theorem 1 are only interesting from an asymptotic point of view. In this section, we show how a reduction from anti-string guessing to paging gives the same asymptotic lower bound but with (much) better constants than those hidden in the $\Omega(n)$ lower bound from Theorem 1. We remark that in the reduction from Anti-q-SG to paging, we need the fact that the lower bound obtained in Theorem 10 is proved using the uniform input distribution.

Theorem 12. *Let $k \in \mathbb{N}$ and let $1 < c < H_k$. For every $\varepsilon > 0$, there exists an n_ε such that a strictly c-competitive paging algorithm with a cache of size k must read at least*

$$b \geq K_{\frac{1}{k+1}}\left(\frac{c}{(k+1)H_k - \varepsilon} + \frac{c}{n}\right) n$$

bits of advice on inputs of length $n \geq n_\varepsilon$.

Proof. We will prove the lower bound by a reduction from Anti-$(k+1)$-SG to paging with a cache of size k and a universe of $k+1$ different pages, $\{s_1, \ldots, s_{k+1}\}$. To this end, let ALG be an arbitrary c-competitive deterministic paging algorithm with advice complexity $b(n)$ and an additive constant of α. Define ALG$'$ to be the following algorithm for Anti-$(k+1)$-SG: For an input string $x = x_1 \ldots x_n \in [k+1]^n$, the algorithm ALG$'$ simulates the paging algorithm ALG on the input $\sigma_x = (s_{x_1}, s_{x_2}, \ldots, s_{x_n})$ with an initial cache of, say, $\{s_1, \ldots, s_k\}$. When ALG$'$ has to output an answer $y_i \in [k+1]$ in round i, it computes the page, s_j, which is outside of the cache of ALG when the ith request of σ_x arrives. ALG$'$ then answers $y_i = j$. Clearly, ALG$'$ incurs a cost of 1 in round i if and only if ALG makes a page-fault on the ith request. Furthermore, ALG$'$ only needs to read $b(n)$ bits of advice for simulating ALG.

Now, we let $p : [k+1]^n \to [0,1]$ be the uniform distribution over all strings of length n (with an alphabet of size $k+1$). We know from Theorem 10 that if $\mathbb{E}_{x\sim p}[\texttt{ALG}'(x)] \leq \beta n$ for every input x of length n, then $b(n) \geq K_{1/q}(\beta)n$. Furthermore, it is known (see e.g. [27]) that for every $\varepsilon > 0$, there exists an n_ε such that $\mathbb{E}_{x\sim p}[\texttt{OPT}(\sigma_x)] \leq \frac{n}{(k+1)H_k - \varepsilon} + 1$ whenever $n \geq n_\varepsilon$. For the rest of the proof, assume that $n \geq n_\varepsilon$. Using that ALG is strictly c-competitive, we have

$$\begin{aligned}\mathbb{E}_{x\sim p}[\texttt{ALG}'(x)] &= \mathbb{E}_{x\sim p}[\texttt{ALG}(\sigma_x)] \\ &\leq c\,\mathbb{E}_{x\sim p}[\texttt{OPT}(\sigma_x)] \\ &\leq \frac{cn}{(k+1)H_k - \varepsilon} + c.\end{aligned}$$

By construction, the number of advice bits read by ALG$'$ is $b(n)$. It follows from Theorem 10 that

$$b(n) \geq K_{\frac{1}{k+1}}\left(\frac{c}{(k+1)H_k - \varepsilon} + \frac{c}{n}\right) n.$$

□

8.4 Weighted binary string guessing

Let $0 < s \leq t$. The *weighted binary string guessing problem* with weights (s,t), denoted BSG-(s,t) for short, is the repeated matrix game with cost-matrix

$$A = \begin{pmatrix} 0 & s \\ t & 0 \end{pmatrix}.$$

For simplicity, we will use $\{0,1\}$ as the alphabet instead of $\{1,2\}$ when working with weighted binary string guessing. It is easy to show that the value of the one-shot game defined by the matrix A is $V = st/(s+t)$, and that the optimal strategy for an algorithm is to output 0 with probability $s/(s+t)$ and 1 with probability $t/(s+t)$. Let $\varepsilon > 0$. From Theorem 3, we get that an online algorithm which on inputs of length n is guaranteed to incur a cost of at most $(st/(s+t) - \varepsilon)\,n$ must read at least $\Omega(n)$ bits of advice. On the other hand, according to Theorem 8, the algorithm that answers 0 with probability $s/(s+t)$ and 1 with probability $t/(s+t)$ incurs an expected cost of at most $(st/(s+t))n$ on every input of length n. Note that for $s = t = 1$, these results are in agreement with the results for (unweighted) binary string guessing.

Theorem 13. *Let $0 < s \leq t$ and let $0 < \alpha < \frac{st}{s+t}$. A randomized BSG-$(s,t)$ algorithm which on every input of length n incurs an expected cost of at most αn must read at least*

$$b \geq K_{s/(s+t)}\left(\frac{\alpha}{t}\right) n$$

bits of advice.

Proof. Let $t > 1$. We start by proving the lower bound for BSG-$(1,t)$. To this end, we use the following n-round input distribution: In each round, the correct answer is 0 with probability $t/(t+1)$ and 1 with probability $1/(t+1)$. Fix a deterministic algorithm ALG which reads at most b bits of advice on on inputs of length n. Let $1 \leq i \leq n$ and let $w \in \mathcal{W}_i$ be any possible history in round i. Recall that $p_i(0|w) = \Pr[X_i = 0|W_i = w]$ and similar for $p_i(1|w)$. Note that if ALG answers 0 in round i, then its expected cost is $p_i(1|w) \cdot t$, and if ALG answers 1, then its expected cost is $p_i(0|w)$. It follows that

$$\mathbb{E}[\text{cost}_i(\texttt{ALG})|W_i = w] \geq \min\{p_i(0|w), p_i(1|w) \cdot t\}. \tag{44}$$

Let $d = D_{KL}(p_i(x|w)\|p_i)$. It follows from (11) of Lemma 1 that $p_i(0|w) \geq K^{-1}_{p(0),l}(d)$ and $p_i(1|w) \geq K^{-1}_{p(1),l}(d)$. A lengthy but straightforward calculation reveals that for all $x \in [0,1]$, it holds that $K_{p(1)}(x/t) \leq K_{p(0)}(x)$. Using this inequality with $x = K^{-1}_{p(0)}(d)$ yields $K^{-1}_{p(1),l}(d) \cdot t \leq K^{-1}_{p(0)}(d)$. Thus, $\min\{p_i(0|w), p_i(1|w) \cdot t\} \geq K^{-1}_{p(1),l}(d) \cdot t$ and therefore

$$\mathbb{E}[\text{cost}_i(\texttt{ALG})|W_i = w] \geq K^{-1}_{p(1),l}(d) \cdot t. \tag{45}$$

Let $f(d) = K^{-1}_{p(1),l}(d) \cdot t$. Since f is convex and decreasing, it follows from Theorem 7 that $\mathbb{E}[\text{cost}(\texttt{ALG})] \geq nf(b/n)$. We will now calculate how large b needs to be in order to ensure that the expected cost is at most αn for $0 < \alpha < t/(1+t)$:

$$\begin{aligned} &\alpha n \geq nf(b/n) \\ &\Rightarrow \alpha \geq K^{-1}_{p(1),l}(b/n)t \\ &\Rightarrow K_{p(1)}(\alpha/t) \leq b/n \\ &\Rightarrow b \geq nK_{1/(1+t)}(\alpha/t). \end{aligned} \tag{46}$$

The general case. We are now ready to prove the lower bound for the general case. Let $0 < \alpha < \frac{st}{s+t}$. Note that in particular, $\alpha < t$. If $s = t$, then BSG-(s,t) is equivalent (after rescaling the costs) to 2-SGKH and we are done. Assume therefore (without loss of generality) that $s < t$. Let ALG be a BSG-(s,t) algorithm with advice complexity $b(n)$. Let ALG$'$ be the BSG-$(1,t/s)$ algorithm which works exactly as ALG. By definition, for every input string $x \in \{0,1\}^n$, it holds that $\texttt{ALG}(x) = s \cdot \texttt{ALG}'(x)$. It follows that if ALG incurs an expected cost of at most αn on inputs of length n, then ALG$'$ never incurs an expected cost larger than $(\alpha/s)n$ for any input string of length n. It then follows from (46) that

$$b(n) \geq K_{s/(s+t)}\left(\frac{\alpha}{t}\right) n$$

□

8.4.1 Proof of Theorem 4: Bin packing

We will show how our lower bound for BSG-(s,t) gives rise to an improved advice complexity lower bound for the classical bin packing problem. This illustrates that even for a problem which is not itself repeatable, it may still be possible to use our techniques to obtain improved lower bounds via e.g. a reduction from a repeated matrix game.

Boyar et al. gave a reduction from binary string guessing (2-SGKH) to bin packing and used this to obtain a lower bound of 9/8 for bin packing algorithms with sublinear advice [31]. Later, Angelopoulos et al. slightly modified the reduction and improved the lower bound to 7/6 [6]. We will show that using the same reduction as in [6], but reducing from *weighted* string guessing further improves the lower bound to $4 - 2\sqrt{2}$. Note that $7/6 = 1.166\overline{6}$ while $4 - 2\sqrt{2} > 1.1715$. We start by describing the reduction from [6].

Reduction from string guessing to bin packing (from [6, 31]). Let $x = x_1 \dots x_n \in \{0,1\}^n$ be a binary string and let $|x|_1$ be the hamming weight of x. We define a bin packing instance, σ_x, as follows: σ_x consists of three phases. Phase 1 consists of $|x|_1$ items all of size $1/2 + \varepsilon$. Phase 2 consists of n items with sizes in $(1/3, 1/2 - \varepsilon)$. Initially, we set $l_0 = 1/3$ and $h_0 = 1/2 - \varepsilon$ and $m_0 = \frac{h_0 + l_0}{2}$. For $1 \leq i \leq n-1$, if $x_i = 0$ then let $l_i = m_{i-1}$ and $h_i = h_{i-1}$, and if $x_i = 1$ then let $l_i = l_{i-1}$ and $h_i = m_{i-1}$. In both cases, let $m_i = (l_i + h_i)/2$. The ith item of phase 2 will have size m_i. If $x_i = 0$, we say that the ith item is *small* and if $x_i = 1$, we say that it is *large*. Note that exactly $|x|_1$ items will be large and $n - |x|_1$ items will be small. Let I_s be the set of small items and let I_l be the set of large items. Phase 3, the final phase of σ_x, consists of $n - |x|_1$ items: For each small item $a \in I_s$, phase 3 contains an item of size $1 - a$ (here, we follow the standard bin packing convention of using the same symbol for both the item and the size of the item).

Note that $|\sigma_x| = 2n$ and $\texttt{OPT}(\sigma_x) = n$. Indeed, an optimal offline algorithm will pack all large items on top of a phase 1 item and will open a new bin for all small items. It will then pack all phase 3 items on top of the corresponding small item.

Given a bin packing algorithm, ALG, we define a string guessing algorithm, ALG′, as follows: On input $x \in \{0,1\}^n$, the algorithm ALG′ will simulate ALG on σ_x. In order to do this, ALG′ is provided with the advice, φ_x, read by ALG on input σ_x. However, note that in order for ALG′ to know how long phase 1 of σ_x is, ALG′ will need to know $|x|_1$. Thus, $|x|_1$ is written onto the advice tape of ALG′ in a self-delimiting way using $2\lceil \log n \rceil + 1 = O(\log n)$ bits. Since ALG′ knows $x_1 \dots x_{i-1}$ when it has to output its guess, y_i, in round i (and since it knows $|x|_1$ and φ_x from its advice), it is clear that ALG′ can determine into which bin ALG packs the ith item, a_i, of phase 2 before deciding on its answer y_i. If a_i is packed on top of a phase 1 item, then ALG′ will answer $y_i = 1$. If a_i is used to open a new bin, then ALG′ will answer $y_i = 0$. If a_i is packed on top of another phase 2 item, then ALG′ will answer $y_i = 1$. Intuitively, ALG′ answers 0 if the packing of a_i indicates that ALG thinks a_i is a small item (and ALG′ answers 1 if ALG seems to think that a_i is large).

We say that a string guessing algorithm makes a 0-*mistake* if it answers 1 when the correct answer is 0, and we define 1-*mistake* symmetrically. Let e_0 be the number of 0-mistakes made by ALG′ and let e_1 be the number of 1-mistakes. Angelopoulos et al. proves in [6] the following relation between the number and type of mistakes made by ALG′ on the input string x and the number of additional (when compared to OPT) bins opened by ALG when given σ_x as input:

$$e_0 \leq 2(\texttt{ALG}(\sigma_x) - n), \tag{47}$$

$$e_1 \leq (\texttt{ALG}(\sigma_x) - n). \tag{48}$$

They then observe that this implies $e_0 + e_1 = \texttt{ALG}'(x) \leq 3(\texttt{ALG}(\sigma_x) - n)$ and use the hardness result for 2-SGKH to obtain a lower bound for bin packing with sublinear advice. Note that by considering the unweighted sum $e_0 + e_1$, the asymmetry present in the reduction between 0-mistakes and 1-mistakes is not being taken into account. We will now show that using our hardness result for weighted string guessing, we can use this asymmetry between 0-mistakes and 1-mistakes in the

reduction to obtain a better lower bound for bin packing. For convenience, we restate Theorem 4 before giving the proof.

Theorem 4. *Let $c < 4 - 2\sqrt{2}$ be a constant. A randomized c-competitive bin packing algorithm must read at least $\Omega(n)$ bits of advice.*

Proof of Theorem 4. Fix $t \geq 1$. The proof is by a reduction from BSG-$(1, t)$. We will show that if

$$c < \frac{t}{(1+t)(2+t)} + 1 \tag{49}$$

and c is constant, then every c-competitive randomized bin packing algorithm must read at least $\Omega(n)$ bits of advice.

Fix $n \in \mathbb{N}$. Let $\texttt{ALG}$ be a deterministic bin packing algorithm reading $b(2n)$ bits of advice on inputs of length $2n$. Using the reduction of Angelopoulos et al. [6] (which is described in details before the statement of Theorem 4), we convert $\texttt{ALG}$ into a BSG-$(1, t)$ algorithm, $\texttt{ALG}'$, which on input $x \in \{0,1\}^n$ works by simulating $\texttt{ALG}$ on the bin packing instance σ_x of length $2n$. From (47) and (48), it follows that $\texttt{ALG}'$ will make at most $2(\texttt{ALG}(\sigma) - n)$ 0-mistakes and at most $(\texttt{ALG}(\sigma) - n)$ 1-mistakes. Thus, $\texttt{ALG}'(x) = e_0 + te_1 \leq (2+t)(\texttt{ALG}(\sigma_x) - n)$.

Let $\texttt{R}$ be a randomized c-competitive bin packing algorithm reading $b(2n)$ bits of advice on inputs of length $2n$. We can convert $\texttt{R}$ into a randomized BSG-$(1, t)$ algorithm, $\texttt{R}'$, by applying the reduction described above to all algorithms in the support of the probability distribution over deterministic algorithms specified by $\texttt{R}$. Doing so, we get that $\mathbb{E}[\texttt{R}'(x)] \leq (2+t)(\mathbb{E}[\texttt{R}(\sigma_x)] - n)$. Since $\texttt{R}$ is c-competitive, this implies that $\mathbb{E}[\texttt{R}'(x)] \leq (2+t)(cn + \alpha - n) = (2+t)(c-1)n + \alpha(2+t)$ for some additive constant α (here, we used that $\texttt{OPT}(\sigma_x) = n$).

Since c is a constant satisfying (49), an easy calculation yields $\mathbb{E}[\texttt{R}'(x)] \leq (t/(1+t) - \varepsilon)n + \alpha(2+t)$ for some constant $\varepsilon > 0$. Since $\alpha(2+t)$ is constant, this implies that for all n large enough, we have $\mathbb{E}[\texttt{R}'(x)] \leq (t/(1+t) - \frac{\varepsilon}{2})n$. Recall that $|\sigma_x| = 2n$ and that $\texttt{R}'$ therefore reads $b(2n) + 2\lceil \log n \rceil + 1$ bits of advice when given x as input (because of how the reduction from string guessing to bin packing works). From Theorem 13 (or simply Theorem 3), we get that $b(2n) + 2\lceil \log n \rceil + 1 = \Omega(n)$. This implies that $b(n) = \Omega(n)$. The proof is finished by setting $t = \sqrt{2}$ for which (49) attains its maximum value of $4 - 2\sqrt{2}$.

□

9 Superlinear advice lower bound for graph coloring

Up until now, we have almost exclusively shown linear advice complexity lower bounds. In this section, we will show to use our direct product theorems to obtain the following superlinear lower bound for online graph coloring (see [71] for a survey on the problem): For every fixed $\varepsilon > 0$, a randomized graph coloring algorithm with advice complexity $o(n \log n)$ must have a competitive ratio of at least $\Omega(n^{1-\varepsilon})$. Previously, it was only known that $\Omega(n \log n)$ bits of advice are were necessary to be 1-competitive [54]. Note that $O(n \log n)$ bits of advice trivially suffice to be optimal. Thus, our lower bound for graph coloring shows that there is a drastic phase-transition in the advice complexity of the problem.

It is easy to see that for every $c = n^{1-o(1)}$, there exists a c-competitive graph coloring algorithm reading $o(n \log n)$ bits of advice. Indeed, let k be the chromatic number of the input graph and suppose that $n/k \leq c$. Then we may just color the graph greedily, since even using n colors is good

enough to be c-competitive. On the other hand, if $n/k > c$, then $k < n/c$. But if $c = n^{1-o(1)}$, then this implies that $k = n^{o(1)}$. Using $b = n\lceil \log k \rceil = n \log(n^{o(1)})$ bits of advice (plus some small lower-order term to encode k in a self-delimiting way) is therefore enough to be c-competitive. But if $b = n \log(n^{o(1)})$, then $b = o(n \log n)$. For example, one can obtain a competitive ratio of $O(n/\log^{O(1)} n)$ using $O(n \log(\log^{O(1)} n)) = O(n \log \log n)$ bits of advice. Thus, our lower bound saying that $\Omega(n \log n)$ bits are needed to be $O(n^{1-\varepsilon})$-competitive for every constant $\varepsilon > 0$ is essentially tight, in the sense that $O(n^{1-\varepsilon})$ cannot be increased to $n^{1-o(1)}$.

In order to prove a lower bound for algorithms with superlinear advice, we apply our martingale-theoretic direct product theorem (Theorem 8) to a hard input distribution for graph coloring due to Halldórsson and Szegedy [61]. They show that a randomized graph coloring algorithm without advice must have a competitive ratio of at least $\Omega(n/\log^2 n)$. By slightly modifying some of the parameters in their construction and combining it with Theorem 8, we obtain a lower bound of $\Omega(n \log n)$ on the advice complexity of algorithms with a competitive ratio of $O(n^{1-\varepsilon})$. We remark that one can equally well use Theorem 7 to obtain this lower bound, but we believe that Theorem 8 is easier to apply in this case.

We begin by formally defining the abstract guessing game which underlies the lower bound in [61].

Definition 13 (cf. [61])**.** The Halldórsson-Szegedy Guessing Game (HSGG) is a minimization problem. At the beginning, two integers k and n where $k \leq n$ and k is even are revealed to the algorithm. The problem consists of n rounds. For $1 \leq i \leq n$, round i proceeds as follows:

1. A set $A_i \in [k]^{k/2}$ of *available characters* is revealed to the algorithm.
2. The algorithm answers $y_i \in [n]$ subject to the following feasibility constraint:

$$\text{For each } 1 \leq t < i, \text{ if } y_t = y_i, \text{ then } x_t \text{ must belong to } A_i. \tag{50}$$

3. The *correct character* $x_i \in A_i$ is revealed to the algorithm.

The cost of the *output* $y = y_1 \dots y_n \in [n]^n$ is the number of different characters from $[n]$ in y, i.e., the cost is $|\{y_1, \dots, y_n\}|$.

Note that there will always be at least one character in $[n]$ satisfying the feasibility constraint (50). This follows since the game has n rounds and since a character which have never been used before by the algorithm is always feasible. Furthermore, we remark that the correct character x_i will not in general be a feasible answer for an online algorithm in round i. Namely, it could be that the algorithm in a previous round $t < i$ (wrongly) answered $y_t = x_i$. If $x_t \notin A_i$, then answering x_i is not possible in round i. However, the cost of an optimal solution to an instance of HSSG is always at most k. Indeed, since `OPT` knows the correct characters $x_1 \dots x_n \in [k]^n$, `OPT` may simply answer $y_i = x_i$ in each round. To see that this is possible, note that if $y_t = y_i$ for some $1 \leq t < i$, then $y_t = x_i$ and hence $x_t = y_t = x_i \in A_i$. Thus, x_i will be a valid answer in round i since `OPT` has not previously made any mistakes.

While k is an upper bound on `OPT`, it is not a tight bound (even if all k characters actually appears in $x_1 \dots x_n$). Consider the instance with $n = 4, k = 4$, correct characters $x_1x_2x_3x_4 = 1234$, and $A_1 = A_2 = \{1, 2\}$, $A_3 = A_4 = \{3, 4\}$. In this case, $y = 1133$ is a feasible output of cost 2.

Note that $O(n \log k)$ bits of advice allows for an optimal HSGG-algorithm.

Theorem 14. *Let $c \in \mathbb{N}$, let $\delta > 1/c$, and let* `ALG` *be a deterministic HSGG-algorithm such that on inputs where $n = k^c$, the cost of* `ALG` *is at most $O(n^{1-\delta})$. Then, the algorithm* `ALG` *must read at least $b = \Omega(n \log n)$ bits of advice.*

Proof (cf. [61]). Fix $c \in \mathbb{N}$. Let k be an even square number and let $n = k^c$. We want to apply Theorem 8. To this end, we construct an n-round input distribution. Each round of the distribution will consist of exactly one request (and hence correspond to a single round of HSGG). In each of the $n = k^c$ rounds, the set of available characters $A_i \in [k]^{k/2}$ is selected uniformly at random (and independently of all previous choices made by the adversary). The correct character $x_i \in A_i$ is then selected uniformly at random from A_i.

Recall that when specifying an n-round input distribution, we also need to specify a cost function for each round. Following [61], we will not use directly the cost function from the definition of HSGG. Instead, for each $1 \leq i \leq n$, we define an auxiliary cost function, pairs_i, as follows: Let $y = y_1 \dots y_n$ be the output computed by some algorithm and let $x = x_1 \dots x_n$ be the correct characters. If $(y_i, x_i) \notin \{(y_t, x_i) : 1 \leq t < i\}$, i.e., round i is the first time the algorithm answers y_i while the correct character is x_i, then we say that (y_i, x_i) is a *new pair* (and that y_i and x_i are *paired*). If (x_i, y_i) is a new pair, then $\text{pairs}_i(y, x) = 1$. Otherwise, if y_i and x_i were already paired previous to round i, then $\text{pairs}_i(y, x) = 0$. It was observed in [61] that if $\text{pairs}(y, x) = \sum_{i=1}^{n} \text{pairs}_i(y, x) = m$, then the actual cost of the output y is at least $\frac{m}{k/2}$ (since $|A_i| = k/2$ for all i, an output character cannot be paired with more than $k/2$ correct characters).

Let $\mathtt{D}$ be a deterministic HSSG-algorithm without advice. We want to show that the probability that $\text{pairs}_i(\mathtt{D})$ is no more than $n/2$ is very small. For a round i and an output character $j \in [n]$, let $H_i(j) = \{x_t \colon y_t = j \text{ and } t < i\}$ be those correct characters which prior to round i were paired with j (note that $H_i(j)$ depends on $\mathtt{D}$). Define

$$\text{dist}(A_i, \mathtt{D}) = \min_{j \in [n] \colon H_i(j) \subseteq A_i} |A_i \setminus H_i(j)| \,. \tag{51}$$

We remark that $\text{dist}(A_i, \mathtt{D})$ is a random variable. Suppose that $\mathtt{D}$ outputs y_i in round i. Then it must hold that $H_i(y_i) \subseteq A_i$ since otherwise y_i was not a feasible answer in that round. Furthermore, since x_i is selected uniformly at random from A_i, the probability that (y_i, x_i) is a new pair is $\frac{|A_i \setminus H_i(y_i)|}{|A_i|}$. Thus, for every history $w \in \mathcal{W}_i$,

$$\Pr\left[\text{pairs}_i(\mathtt{D}) = 1 | W_i = w\right] \geq \frac{\mathbb{E}[\text{dist}(A_i, \mathtt{D}) | W_i = w]}{k/2}. \tag{52}$$

The random variable $\text{dist}(A_i, \mathtt{D})$ depends in a complicated way on the previous rounds of the guessing game. However, suppose we want an upper bound on $\Pr[\text{dist}(A_i, \mathtt{D}) \leq s]$ for some $s \in \mathbb{N}$. No matter how the previous rounds have proceeded, for each of the n output characters $j \in [n]$, the set $H_i(j)$ is a subset of $[k]$ of size at most $k/2$. Furthermore, $H_i(j)$ is a subset of at most $\binom{k-|H_i(j)|}{k/2-|H_i(j)|}$ sets of size $k/2$. Thus, if $|H_i(j)| \geq k/2 - s$, then $H_i(j)$ is a subset of at most $\binom{k/2+s}{s}$ sets of size $k/2$. Using that there are $\binom{k}{k/2}$ choices for A_i (each of which is equally likely) and the

union bound, we get that

$$\begin{aligned}\Pr\left[\operatorname{dist}(A_i,\mathtt{D}) \le k/2-\sqrt{k}\middle| W_i = w\right] &\le \frac{\binom{k-\sqrt{k}}{k/2-\sqrt{k}}\cdot n}{\binom{k}{k/2}}\\ &\le \frac{2^{k-\sqrt{k}}\cdot k^c}{2^k\cdot k^{-1}}\\ &= 2^{-\sqrt{k}}k^{c+1}.\end{aligned}$$

Here, we used that $\binom{k}{k/2} \ge 2^k/k$. We can now bound from below the expected value of $\operatorname{dist}(A_i,\mathtt{D})$. Set $s = k/2-\sqrt{k}$. Then,

$$\begin{aligned}\mathbb{E}\left[\operatorname{dist}(A_i,\mathtt{D})|W_i=w\right] &\ge s\cdot\Pr\left[\operatorname{dist}(A_i,\mathtt{D})\ge s|W_i=w\right]\\ &\ge s\cdot\left(1-\Pr\left[\operatorname{dist}(A_i,\mathtt{D})\le s|W_i=w\right]\right)\\ &\ge \left(k/2-\sqrt{k}\right)\left(1-2^{-\sqrt{k}}k^{c+1}\right).\end{aligned}$$

Thus,

$$\begin{aligned}\frac{\mathbb{E}\left[\operatorname{dist}(A_i,\mathtt{D})|W_i=w\right]}{k/2} &\ge \frac{\left(k/2-\sqrt{k}\right)\left(1-2^{-\sqrt{k}}k^{c+1}\right)}{k/2}\\ &\ge \left(1-\frac{2}{\sqrt{k}}\right)\cdot\left(1-\frac{k^{c+1}}{2^{\sqrt{k}}}\right)\\ &\ge \left(1-\frac{2}{\sqrt{k}}\right)^2 \text{ for } k\ge K.\end{aligned}$$

Here, K is chosen such that $(1-2/\sqrt{k}) \le (1-2^{-\sqrt{k}}k^{c+1})$ for all $k \ge K$. This is always possible since $k^{c+1}\sqrt{k} = o(2^{\sqrt{k}})$. Set $p := (1-2/\sqrt{k})^2$. For the rest of the proof, we assume $k \ge K$, and we assume $k \ge 49$ so that $p \ge 1/2$.

Just as $\operatorname{dist}(A_i,\mathtt{D})$, the random variable $\operatorname{pairs}_i(\mathtt{D})$ is obviously not independent of $\operatorname{pairs}_t(\mathtt{D})$ for $t < i$. However, the lower bound derived above holds for every possible history $w \in \mathcal{W}_i$ of the game at round i. Thus, $\operatorname{pairs}_i(\mathtt{D})$ is a $\{0,1\}$-valued random variable with conditional expectation $\mathbb{E}[\operatorname{pairs}_i(\mathtt{D})|W_i=w] = \Pr[\operatorname{pairs}_i(\mathtt{D})=1|W_i=w] \ge p$ for every $w \in \mathcal{W}_i$. From this, it follows that for every $w \in \mathcal{W}_i$,

$$\begin{aligned}\mathbb{E}\left[(\operatorname{pairs}_i(\mathtt{D})-p)^2\middle|W_i=w\right] &= \Pr[\operatorname{pairs}_i(\mathtt{D})=0|W_i=w]\cdot(0-p)^2\\ &\quad+\Pr[\operatorname{pairs}_i(\mathtt{D})=1|W_i=w]\cdot(1-p)^2\\ &= (1-\Pr[\operatorname{pairs}_i(\mathtt{D})=1|W_i=w])\cdot p^2\\ &\quad+\Pr[\operatorname{pairs}_i(\mathtt{D})=1|W_i=w]\cdot(1-p)^2\\ &\le p(1-p).\end{aligned}$$

The last inequality above holds since we are assuming that $p \ge 1/2$. let ALG be a deterministic HSGG-algorithm with advice complexity $b(n)$ such that on inputs where $n = k^c$, the cost of ALG is

at most $O(n^{1-\delta})$. If we let $\varepsilon = p - 1/2$, it follows from Theorem 8 and our calculations above that

$$\begin{aligned}\Pr[\text{pairs}(\texttt{ALG}) \leq 1/2n] &= \Pr[\text{pairs}(\texttt{ALG}) \leq (p-\varepsilon)n] \\ &\leq \exp_2\left(b(n) - K_{\frac{\gamma}{1+\gamma}}\left(\frac{\alpha+\gamma}{1+\gamma}\right)n\right) \\ &= \exp_2\left(b(n) - K_{1-p}\left(\frac{1}{2}\right)n\right),\end{aligned}$$

where $\gamma = (p(1-p))/p^2 = \frac{1-p}{p}$ and $\alpha = \varepsilon/p = 1 - \frac{1}{2p}$. Now,

$$\begin{aligned}K_{1-p}(1/2) &= 1/2 \cdot \log\left(\frac{1/2}{1-p}\right) + 1/2 \cdot \log\left(\frac{1/2}{p}\right) \\ &= 1/2 \cdot \log\left(\frac{1}{4p(1-p)}\right) \\ &= 1/2 \cdot \log\left(\frac{k^2}{16(\sqrt{k}-2)^2(\sqrt{k}-1)}\right) \\ &\geq 1/2 \cdot \log\left(\frac{k^2}{16(\sqrt{k})^2(\sqrt{k})}\right) \\ &= 1/2 \cdot \log\left(\frac{\sqrt{k}}{16}\right).\end{aligned}$$

Thus,

$$\Pr[\text{pairs}(\texttt{ALG}) \leq n/2] \leq \exp_2\left(b(n) - \frac{1}{2}\log\left(\frac{\sqrt{k}}{16}\right)n\right). \tag{53}$$

Recall that the cost of ALG is assumed to be at most $O(n^{1-\delta})$ for some fixed constant $\delta > 1/c$. Assume by way of contradiction that $b(n) \in o(n \log n)$. Under this assumption, if k (and thereby also $n = k^c$) is sufficiently large, then $b(n) < \frac{1}{2}\log\left(\frac{\sqrt{k}}{16}\right)n$ since $\frac{1}{2}\log\left(\frac{\sqrt{k}}{16}\right)n = \frac{1}{2}\log\left(\frac{n^{1/c}}{16}\right)n = \Omega(n \log n)$. Thus, for all sufficiently large k and $n = k^c$, it follows from (53) that there exists an input, σ, such that $\text{pairs}(\texttt{ALG}(\sigma)) \geq n/2$ and, hence, $\texttt{ALG}(\sigma) \geq \frac{\text{pairs}(\texttt{ALG}(\sigma))}{k/2} \geq \frac{n/2}{k/2} = n^{1-1/c}$. This contradicts the claimed performance guarantee of ALG since $1 - 1/c > 1 - \delta$. We conclude that $b(n) \in \Omega(n \log n)$.

□

The following lemma which provides a reduction from HSGG to graph coloring follows directly from [61].

Lemma 11 ([61])**.** *If there exists a deterministic graph-coloring algorithm with advice complexity $b(n)$ which for n-vertex k-colorable graphs uses at most $f(n,k)$ colors, then there exists a HSGG-algorithm with advice complexity $b(n)$ which always incurs a cost of at most $f(n,k)$.*

Combining Theorem 14 with Lemma 11, we obtain a proof of Theorem 5 (we restate the theorem here for convenience).

Theorem 5. *Let $\varepsilon > 0$ be a constant. A randomized $O(n^{1-\varepsilon})$-competitive online graph coloring algorithm must read at least $\Omega(n \log n)$ bits of advice.*

Proof of Theorem 5. Let ALG be a deterministic $c(n)$-competitive graph coloring algorithm where $c(n) \in O(n^{1-\varepsilon})$. By definition, $\texttt{ALG}(\sigma) \leq c(n) \cdot \texttt{OPT}(\sigma) + \alpha$ for some constant α. In particular, this means that ALG can color n-vertex k-colorable graphs online using at most $f(n,k) = c(n) \cdot k + \alpha$ colors. By the reduction from HSGG to graph coloring (Lemma 11), we get a HSGG-algorithm, ALG′, which incurs a cost of at most $f(n,k)$ and which has the same advice complexity as ALG.

Choose $c \in \mathbb{N}$ such that $2/c < \varepsilon$ and consider HSGG-inputs where $n = k^c$. On these inputs, ALG′ incurs a cost of at most $f(n, n^{1/c}) = c(n) \cdot n^{1/c} + \alpha = O(n^{1-\varepsilon} n^{1/c}) = O(n^{1-\delta})$ where $\delta := \varepsilon - 1/c$. Since $2/c < \varepsilon$, it follows that $\delta = \varepsilon - 1/c > 1/c$. By Theorem 14, ALG′ must therefore use $\Omega(n \log n)$ bits of advice. Since the advice complexity of ALG′ was the same as that of ALG, this proves Theorem 5 for deterministic algorithms.

We still need to consider the case where the graph coloring algorithm is randomized. Note that for a graph on n vertices, the number of requests in a single round can be at most 2^n (there are at most 2^n possible subsets of incident vertices for the vertex currently being revealed). Thus, the number of inputs of length n for online graph coloring is at most $(2^n)^n = 2^{n^2} = 2^{n^{O(1)}}$. It follows from Theorem 16 in Appendix A that a randomized $O(n^{1-\varepsilon})$-competitive graph coloring algorithm with advice complexity $b(n)$ can be converted to a deterministic $O(n^{1-\varepsilon})$-competitive graph coloring algorithm with advice complexity $b(n) + O(\log n)$. From the above, we get that $b(n) + O(\log n) = \Omega(n \log n)$ from which it follows that $b(n) = \Omega(n \log n)$. □

10 Lower bounds for ∨-repeatable problems

So far, we have only considered Σ-repeatable (minimization) problems, where the cost of a solution in P^* is the sum of costs in each round. In this section, we consider ∨-repeatable (minimization) problems, where the cost of a solution in P^* is the *maximum* of the costs in each round. Theorem 15 shows that lower bounds for ∨-repeatable problems can often be obtained rather easily. We remark that the proof of Theorem 15 does not rely on our previous direct product theorems.

Theorem 15. *Let* P *be a strictly ∨-repeatable problem and let $I = \{\sigma_1, \ldots, \sigma_m\}$ be a finite set of* P*-inputs. Furthermore, let $t = \max_{\sigma \in I} \texttt{OPT}(\sigma)$ and let $\varepsilon > 0$ be a constant. Suppose that for every deterministic* P*-algorithm without advice, ALG, there exists some $1 \leq i \leq m$ such that $\texttt{ALG}(\sigma_i) \geq k$. Then, for every randomized* P*-algorithm, R, reading $o(n)$ bits of advice, there exists a* P*-input σ such that $\mathbb{E}[\texttt{R}(\sigma)] \geq (1-\varepsilon)k$ and such that $\texttt{OPT}(\sigma) \leq t$.*

Proof. The proof is via Yao's principle. Thus, let ALG be an arbitrary deterministic P-algorithm reading $b(n) \in o(n)$ bits of advice on inputs of length n. We will first construct a hard input distribution for $\mathrm{P}^*_\vee$, and then use that P is ∨-repeatable.

For each $r \in \mathbb{N}$, we define an r-round input distribution, p^r, for $\mathrm{P}^*_\vee$ as follows: In each round, the adversary selects uniformly at random an input from I. When a round ends, a reset occurs (thus, each round of the r-round input distribution corresponds to a round of $\mathrm{P}^*_\vee$). Let $n' = \max_{\sigma \in I} |\sigma|$ be the length of the longest input in I. Note that an input in the support of p^r has length at most $n'r$.

By assumption, $b : \mathbb{N} \to \mathbb{N}$ is a function such that $b(n) \in o(n)$. Since n' is a constant, this means that $b(n'r) \in o(r)$. Thus, since $1 - 1/m < 1$, it follows that $2^{b(n'r)}(1-1/m)^r \to 0$ as $r \to \infty$. Choose r large enough that $2^{b(n'r)}(1-1/m)^r \leq \varepsilon$.

The number r of rounds is now fixed. Following Lemma 3, we will define a $\mathrm{P}^*_\vee$-algorithm $\mathtt{ALG}^*$ based on the P-algorithm $\mathtt{ALG}$. The algorithm $\mathtt{ALG}^*$ will only be defined for inputs in $\mathrm{supp}(p^r)$. For other inputs, it can behave arbitrarily. By definition, for each input length n, there exists $2^{b(n)}$ deterministic P-algorithms, $\mathtt{ALG}_1, \ldots, \mathtt{ALG}_{2^{b(n)}}$ such that $\mathtt{ALG}(\sigma) = \min_j \mathtt{ALG}_j(\sigma)$ for every P-input σ of length at most n. Since P is strictly $\vee$-repeatable, this implies that for every n, there exists $2^{b(n)}$ deterministic $\mathrm{P}^*_\vee$-algorithms, $\mathtt{ALG}^*_1, \ldots, \mathtt{ALG}^*_{2^{b(n)}}$ such that $\mathtt{ALG}^*_j(\sigma^*) \leq \mathtt{ALG}_j(g(\sigma^*))$ for every $\mathrm{P}^*_\vee$-input σ^* of length at most n (where g is as in Definition 6). Recall that an input in $\mathrm{supp}(p^r)$ has length at most $n'r$. We define $\mathtt{ALG}^*$ to be the following $\mathrm{P}^*_\vee$-algorithm: For an input $\sigma^* \in \mathrm{supp}(p^r)$, the algorithm $\mathtt{ALG}^*$ will read $b(n'r)$ bits of advice and use these to select the best of the $2^{b(n'r)}$ algorithms $\mathtt{ALG}^*_1, \ldots, \mathtt{ALG}^*_{2^{b(n'r)}}$ for this particular input σ^*. Note that since r is fixed and n' is a constant, $\mathtt{ALG}^*$ can compute $n'r$ itself (and thereby also $b(n'r)$). By definition, $\mathtt{ALG}^*(\sigma^*) \leq \mathtt{ALG}(g(\sigma^*))$ for every input $\sigma^* \in \mathrm{supp}(p^r)$.

Let $1 \leq j \leq 2^{b(n'r)}$ and let $1 \leq i \leq r$. We consider the computation of $\mathtt{ALG}^*_j$ in round i of p^r. The computation of $\mathtt{ALG}^*_j$ in round i may depend on the request revealed in previous rounds. However, recall that a reset occurred just before round i began. Thus, for each fixed history $w \in \mathcal{W}_i$, the computation of the algorithm $\mathtt{ALG}^*_j$ in round i given the history w defines a deterministic P-algorithm valid for all inputs in I. Thus, by assumption, there exists at least one input $\sigma \in I$ such that if σ is the input in round i, then $\mathtt{ALG}^*_j$ incurs a cost of at least k in round i. Since each of the m inputs in I have probability $1/m$ of being selected in round i, this implies that for *every* $w \in \mathcal{W}_i$,

$$\Pr[\mathrm{cost}_i(\mathtt{ALG}^*_j) < k | W_i = w] \leq 1 - 1/m. \tag{54}$$

Thus, even though the costs incurred by $\mathtt{ALG}^*_j$ in each of the r rounds are not (necessarily) independent, we can use (54) to bound from above the probability that the maximum cost over all r rounds is less than k:

$$\begin{aligned}\Pr[\mathrm{cost}(\mathtt{ALG}^*_j) < k] &= \Pr[\forall 1 \leq i \leq r\colon\ \mathrm{cost}_i(\mathtt{ALG}^*_j) < k] \\ &\leq (1 - 1/m)^r.\end{aligned}$$

Applying the union bound over all $2^{b(n'r)}$ algorithms $\mathtt{ALG}^*_1, \ldots, \mathtt{ALG}_{2^{b(n'r)}}$ gives

$$\begin{aligned}\Pr\left[\mathrm{cost}(\mathtt{ALG}^*) < k\right] &= \Pr\left[\exists 1 \leq j \leq 2^{b(n'r)}\colon\ \mathrm{cost}(\mathtt{ALG}^*_j) < k\right] \\ &\leq 2^{b(n'r)}(1 - 1/m)^r \\ &\leq \varepsilon.\end{aligned} \tag{55}$$

It follows that

$$\begin{aligned}\mathbb{E}[\mathrm{cost}(\mathtt{ALG}^*)] &\geq \Pr[\mathrm{cost}(\mathtt{ALG}^*) \geq k] \cdot k \\ &\geq (1 - \varepsilon) \cdot k.\end{aligned}$$

This implies that $\mathbb{E}_{\sigma^* \sim p^r}[\mathtt{ALG}(g(\sigma^*))] \geq \mathbb{E}_{\sigma^* \sim p^r}[\mathtt{ALG}^*(\sigma^*)] \geq (1 - \varepsilon)k$.

Using Yao's principle, we conclude that for every randomized P-algorithm $\mathtt{R}$ with advice complexity $o(n)$, there exists an $r \in \mathbb{N}$ and a $\sigma^* \in \mathrm{supp}(p^r)$ such that $\mathbb{E}[\mathtt{R}(g(\sigma^*))] \geq (1 - \varepsilon)k$. Since $\mathtt{OPT}(g(\sigma^*)) \leq \mathtt{OPT}^*_\vee(\sigma^*) \leq t$ for every $\sigma^* \in \mathrm{supp}(p^r)$, this finishes the proof. □

10.1 Applications

Application: Edge coloring It is easy to see that online edge coloring is strictly $\vee$-repeatable (see [14] for a formal definition of the problem). In [14], it is shown that for every $\Delta \in \mathbb{N}$, there exists a finite set of inputs I_Δ with the following properties: For $\sigma \in I_\Delta$, the underlying graph of σ is a forest (and hence Δ-edge-colorable). Also, for every deterministic algorithm $\texttt{D}$ without advice, there exists some $\sigma \in I_\Delta$ such that $\texttt{D}(\sigma) \geq (2\Delta - 1)$. We can use this construction together with Theorem 15. Let $\varepsilon > 0$ be a fixed constant. Suppose by way of contradiction that $\texttt{ALG}$ is a $(2-\varepsilon)$-competitive randomized edge coloring algorithm with advice complexity $o(n)$. Then there exists some α such that $\mathbb{E}[\texttt{ALG}(\sigma)] \leq (2-\varepsilon)\texttt{OPT}(\sigma) + \alpha$ for all inputs σ. Choose Δ large enough so that $(2-\varepsilon)\Delta + \alpha < 2\Delta - 1$. By Theorem 15, there exists an input $\sigma \in I_\Delta$ such that $\mathbb{E}[\texttt{ALG}(\sigma)] \geq 2\Delta - 1 > (2-\varepsilon)\Delta + \alpha = (2-\varepsilon)\texttt{OPT}(\sigma) + \alpha$. This is a contradiction. Since ε was arbitrary, we conclude that the competitive ratio of a randomized edge-coloring algorithm with advice complexity $o(n)$ must be at least 2. This result was already obtained in [82] where it shown directly how an adversary may repeatedly use the lower bound for algorithms without advice. Theorem 15 reveals that the proof of the lower bound in [82] is in fact just a special case of a much more general technique.

Application: Graph coloring We have already shown previously how to obtain good advice complexity lower bounds for graph coloring on general graphs (see Section 9). We will now consider the 2-vertex-coloring problem (i.e., the online graph coloring problem restricted to bipartite graphs).

It is easy to see that online 2-vertex-coloring is strictly $\vee$-repeatable. Furthermore, it is shown in [60] that for every n which is a power of two, there exists a finite set of inputs, I_n, with the following properties: For $\sigma \in I_n$, the graph defined by the input σ is a 2-colorable graph on n vertices. Furthermore, for every deterministic online graph coloring algorithm $\texttt{D}$ there exists a $\sigma \in I_n$ such that $\texttt{D}(\sigma) \geq \log_2(n)$.

We will now apply Theorem 15. Let $c = O(1)$ be a constant. Assume by way of contradiction that $\texttt{ALG}$ is a c-competitive randomized 2-vertex-coloring algorithm and that the advice complexity of $\texttt{ALG}$ is $o(n)$. This means that there exists an additive constant α such that $\mathbb{E}[\texttt{ALG}(\sigma)] \leq c \cdot \texttt{OPT}(\sigma) + \alpha$ for every input σ where the underlying graph is bipartite. Choose n large enough so that $2c + \alpha < \log_2(n)$ (and such that n is a power of two). By Theorem 15, we get that for this choice of n, there exists an input $\sigma \in I_n$ such that $\mathbb{E}[\texttt{ALG}(\sigma)] \geq \log_2(n) > 2c + \alpha = c \cdot \texttt{OPT}(\sigma) + \alpha$. But this is a contradiction since $\texttt{ALG}$ was assumed to be c-competitive with an additive constant of α. We conclude that if a randomized algorithm with advice complexity $o(n)$ is c-competitive, then it must be the case that $c = \omega(1)$.

Previously, is was only known that $\Omega(n)$ bits of advice was needed to never use more than 3 colors when coloring a bipartite graph online [18]. Thus, it was known that an algorithm with advice complexity $o(n)$ would use 4 colors on some graphs and therefore have a strict competitive ratio of at least 2.

Application: $L(2,1)$-coloring on paths and cycles See [19] for a formal definition of the problem. The input is a graph of maximum degree two. When a vertex is revealed, an online algorithm must irrevocably assign a non-negative integer to the vertex such that the integers assigned to the set of vertices is a valid $L(2,1)$-coloring. The cost of a solution is λ, where λ is the smallest integer such that the set of colors used in the coloring is contained in $\{0, 1, \ldots, \lambda\}$. It is easy to see that $\texttt{OPT}(\sigma) \leq 4$ for every input σ.

In [19], the following results were obtained: It is shown that there exists a deterministic algorithm $\mathtt{D}$ without advice such that $\mathtt{D}(\sigma) \leq 6$ for every input σ. A matching lower bound is provided by showing that there exists a finite set of inputs I such that, for each deterministic algorithm $\mathtt{ALG}$ without advice, $\mathtt{ALG}(\sigma) \geq 6$ for some $\sigma \in I$. Finally, it is also shown in [19] that a deterministic algorithm $\mathtt{ALG}$ which can guarantee to solve the $L(2,1)$ problem on paths with a cost of at most 5 must read at least $3.9402 \cdot 10^{-10} n = \Omega(n)$ bits of advice. Since costs are integral and $\mathtt{OPT}(\sigma)$ is always at most 4, this gives a (tight) lower bound of $3/2$ on the strict competitive ratio of algorithms with sublinear advice. Using derandomization (see Appendix A), the authors conclude that no randomized algorithm can achieve a strict competitive ratio smaller than $3/2$.

It is easy to see that $L(2,1)$-coloring on paths is a strictly $\vee$-repeatable problem. Thus, combining the lower bound for deterministic algorithms without advice from [19] with Theorem 15, we get that for every $\varepsilon > 0$ and every (possibly randomized) algorithm reading $o(n)$ bits of advice, there exists an input σ such that $\mathbb{E}[\mathtt{ALG}(\sigma)] \geq (1-\varepsilon)6$. Since costs are integral, this reproves the lower bound of $3/2$ on the strict competitive ratio of algorithms with $o(n)$ bits of advice. We believe that the proof is simpler than the existing proof. Furthermore, it highlights the fact that obtaining a lower bound for deterministic algorithms without advice suffices in order to obtain a lower bound for randomized algorithms.

Necessity of assumptions in Theorem 15 Note that in order to successfully apply Theorem 15, it must be the case that $t = \max_{\sigma \in I} \mathtt{OPT}(\sigma)$ is small compared to k (where I and k are as in the statement of Theorem 15). One might wonder if it is possible to get rid of this assumption. In particular, one could hope to show that if algorithms without advice for a $\vee$-repeatable problem can be no better than c-competitive, then algorithms with advice complexity $o(n)$ also cannot be (much) better than c-competitive. Unfortunately, it turns out that this is false in general. For instance, the online multi-coloring problem on a path network (without deletions) provides a counter-example (see [33] for a definition of the problem). This problem is strictly $\vee$-repeatable. It is known that no algorithm without advice (even if we allow randomization) can achieve a competitive ratio better than $4/3$ [36]. On the other hand, it has been shown that there exists a strictly 1-competitive algorithm which uses $O(\log n)$ bits of advice [37]. The advice given to the algorithm is the number of colors used in an optimal solution.

The structure of the 4/3-lower bound for randomized algorithms is as follows: For an input length n, one defines two inputs, σ_1^n, σ_2^n, and considers the uniform distribution p over these two inputs. The inputs σ_1^n and σ_2^n are such that $\mathtt{OPT}(\sigma_1) = n/2$ while $\mathtt{OPT}(\sigma_2) = n/4$. This implies that $\mathbb{E}_{\sigma \sim p}[\mathtt{OPT}(\sigma)] = 1/2 \cdot (n/2 + n/4) = (3/8)n$. On the other hand, it is shown that for every deterministic algorithm $\mathtt{D}$ (without advice), it holds that $\mathbb{E}_{\sigma \sim p}[\mathtt{D}(\sigma)] \geq n/2$, which yields a lower bound of $4/3$. However, this lower bound does not carry over to algorithms with advice complexity $o(n)$ by simply repeating it r times as in the proof of Theorem 15. The problem is that the expected cost of $\mathtt{OPT}$ will tend to $n/2$ as r tends to infinity. Indeed, the probability that $\mathtt{OPT}$ incurs a cost of $n/2$ in any given round is $1/2$, and if the cost of $\mathtt{OPT}$ in even a single round is $n/2$ then the total cost of $\mathtt{OPT}$ will also be $n/2$. This example illustrates why we need to work with $t = \max_{\sigma \in I} \mathtt{OPT}(\sigma)$ in Theorem 15.

Acknowledgment. The author would like to thank Joan Boyar, Magnus Find, Lene Favrholdt, Adi Rosén, and the anonymous reviewers for helpful comments on this work and its presentation.

References

[1] Anna Adamaszek, Artur Czumaj, Matthias Englert, and Harald Räcke. Almost tight bounds for reordering buffer management. In *STOC*, pages 607–616. ACM, 2011. doi:10.1145/1993636.1993717.

[2] Anna Adamaszek, Marc P. Renault, Adi Rosén, and Rob van Stee. Reordering buffer management with advice. In *WAOA*, pages 132–143. Springer, 2013. doi:10.1007/978-3-319-08001-7_12.

[3] Susanne Albers and Matthias Hellwig. Online makespan minimization with parallel schedules. In *SWAT*, pages 13–25, 2014. doi:10.1007/978-3-319-08404-6_2.

[4] Susanne Albers, Bernhard von Stengel, and Ralph Werchner. A combined BIT and TIMESTAMP algorithm for the list update problem. *Inf. Process. Lett.*, 56(3):135–139, 1995. doi:10.1016/0020-0190(95)00142-Y.

[5] Christoph Ambühl. *On the list update problem*. PhD thesis, ETH Zürich, 2002. doi:10.3929/ethz-a-004394113.

[6] Spyros Angelopoulos, Christoph Dürr, Shahin Kamali, Marc P. Renault, and Adi Rosén. Online bin packing with advice of small size. In *WADS*, volume 9214 of *Lecture Notes in Computer Science*, pages 40–53. Springer, 2015. doi:10.1007/978-3-319-21840-3_4.

[7] Robert J. Aumann. Mixed and behavior strategies in infinite extensive games. In *Advances in Game Theory*, pages 627–650. Princeton Univ. Press, 1964.

[8] Noa Avigdor-Elgrabli and Yuval Rabani. An optimal randomized online algorithm for reordering buffer management. In *FOCS*, pages 1–10. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.9.

[9] Yossi Azar, Andrei Z. Broder, and Mark S. Manasse. On-line choice of on-line algorithms. In Vijaya Ramachandran, editor, *SODA*, pages 432–440. ACM/SIAM, 1993. URL: http://dl.acm.org/citation.cfm?id=313559.313847.

[10] Yossi Azar, Iftah Gamzu, and Ran Roth. Submodular max-sat. In *ESA*, volume 6942 of *Lecture Notes in Computer Science*, pages 323–334. Springer, 2011. doi:10.1007/978-3-642-23719-5_28.

[11] Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Math. J. (2)*, 19:357–367, 1967.

[12] Nikhil Bansal, Niv Buchbinder, Anupam Gupta, and Joseph Naor. A randomized o(log2 k)-competitive algorithm for metric bipartite matching. *Algorithmica*, 68(2):390–403, 2014. doi:10.1007/s00453-012-9676-9.

[13] Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. A polylogarithmic-competitive algorithm for the k-server problem. *J. ACM*, 62(5):40, 2015. Preliminary version in FOCS'11. doi:10.1145/2783434.

[14] Amotz Bar-Noy, Rajeev Motwani, and Joseph (Seffi) Naor. The greedy algorithm is optimal for on-line edge coloring. *Inf. Process. Lett.*, 44(5):251–253, 1992. doi:10.1016/0020-0190(92)90209-E.

[15] Yair Bartal, Béla Bollobás, and Manor Mendel. Ramsey-type theorems for metric spaces with applications to online problems. *J. Comput. Syst. Sci.*, 72(5):890–921, 2006. doi:10.1016/j.jcss.2005.05.008.

[16] Dwight R. Bean. Effective coloration. *J. Symb. Log.*, 41(2):469–480, 1976. doi:10.1017/S0022481200051549.

[17] Shai Ben-David, Allan Borodin, Richard M. Karp, Gábor Tardos, and Avi Wigderson. On the power of randomization in on-line algorithms. *Algorithmica*, 11(1):2–14, 1994. Preliminary version in STOC'90. doi:10.1007/BF01294260.

[18] Maria Paola Bianchi, Hans-Joachim Böckenhauer, Juraj Hromkovič, and Lucia Keller. Online coloring of bipartite graphs with and without advice. *Algorithmica*, 70(1):92–111, 2014. doi:10.1007/s00453-013-9819-7.

[19] Maria Paola Bianchi, Hans-Joachim Böckenhauer, Juraj Hromkovič, Sacha Krug, and Björn Steffen. On the advice complexity of the online $L(2,1)$-coloring problem on paths and cycles. *Theor. Comput. Sci.*, 554:22–39, 2014. doi:10.1016/j.tcs.2014.06.027.

[20] Avrim Blum and Carl Burch. On-line learning and the metrical task system problem. *Machine Learning*, 39(1):35–58, 2000. Preliminary version in COLT'97. doi:10.1023/A:1007621832648.

[21] Hans-Joachim Böckenhauer, Richard Dobson, Sacha Krug, and Kathleen Steinhöfel. On energy-efficient computations with advice. In *COCOON*, pages 747–758. Springer, 2015. doi:10.1007/978-3-319-21398-9_58.

[22] Hans-Joachim Böckenhauer, Sascha Geulen, Dennis Komm, and Walter Unger. Constructing randomized online algorithms from algorithms with advice. *ETH-Zürich*, 2015. doi:10.3929/ethz-a-010556417.

[23] Hans-Joachim Böckenhauer, Juraj Hromkovič, and Dennis Komm. A technique to obtain hardness results for randomized online algorithms - A survey. In *Computing with New Resources - Essays Dedicated to Jozef Gruska on the Occasion of His 80th Birthday*, volume 8808 of *LNCS*, pages 264–276. Springer, 2014. URL: http://dx.doi.org/10.1007/978-3-319-13350-8_20, doi:10.1007/978-3-319-13350-8_20.

[24] Hans-Joachim Böckenhauer, Juraj Hromkovič, Dennis Komm, Sacha Krug, Jasmin Smula, and Andreas Sprock. The string guessing problem as a method to prove lower bounds on the advice complexity. *Theor. Comput. Sci.*, 554:95–108, 2014. Preliminary version in COCOON'13. doi:10.1016/j.tcs.2014.06.006.

[25] Hans-Joachim Böckenhauer, Dennis Komm, Rastislav Královič, and Richard Královič. On the advice complexity of the k-server problem. In *ICALP (1)*, pages 207–218, 2011. doi:10.1007/978-3-642-22006-7_18.

[26] Hans-Joachim Böckenhauer, Dennis Komm, Rastislav Královič, Richard Královič, and Tobias Mömke. On the advice complexity of online problems. In *ISAAC*, pages 331–340, 2009. doi:10.1007/978-3-642-10631-6_35.

[27] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[28] Allan Borodin, Nathan Linial, and Michael E. Saks. An optimal on-line algorithm for metrical task system. *J. ACM*, 39(4):745–763, 1992. Preliminary version in STOC'87. doi:10.1145/146585.146588.

[29] Joan Boyar, Lene M. Favrholdt, Christian Kudahl, and Jesper W. Mikkelsen. Advice complexity for a class of online problems. In *STACS*, volume 30 of *LIPIcs*, pages 116–129. Schloss Dagstuhl, 2015. doi:10.4230/LIPIcs.STACS.2015.116.

[30] Joan Boyar, Shahin Kamali, Kim S. Larsen, and Alejandro López-Ortiz. On the list update problem with advice. In *LATA*, pages 210–221, 2014. doi:10.1007/978-3-319-04921-2_17.

[31] Joan Boyar, Shahin Kamali, Kim S. Larsen, and Alejandro López-Ortiz. Online bin packing with advice. In *STACS*, pages 174–186, 2014. Full paper to appear in *Algorithmica*. doi:10.4230/LIPIcs.STACS.2014.174.

[32] Carl Burch. *Machine learning in metrical task systems and other on-line problems*. PhD thesis, Carnegie Mellon University, 2000. http://cburch.com/pub/thesis.ps.gz.

[33] Joseph Wun-Tat Chan, Francis Y. L. Chin, Deshi Ye, Yong Zhang, and Hong Zhu. Frequency allocation problems for linear cellular networks. In *ISAAC*, volume 4288, pages 61–70. Springer, 2006. doi:10.1007/11940128_8.

[34] Timothy M. Chan and Hamid Zarrabi-Zadeh. A randomized algorithm for online unit clustering. *Theory Comput. Syst.*, 45(3):486–496, 2009. doi:10.1007/s00224-007-9085-7.

[35] Arkadev Chattopadhyay, Jeff Edmonds, Faith Ellen, and Toniann Pitassi. A little advice can be very helpful. In *SODA*, pages 615–625. SIAM, 2012.

[36] Marie G. Christ, Lene M. Favrholdt, and Kim S. Larsen. Online multi-coloring on the path revisited. *Acta Inf.*, 50(5-6):343–357, 2013. doi:10.1007/s00236-013-0184-4.

[37] Marie G. Christ, Lene M. Favrholdt, and Kim S. Larsen. Online multi-coloring with advice. *Theor. Comput. Sci.*, 596:79–91, 2015. doi:10.1016/j.tcs.2015.06.044.

[38] Marek Chrobak and Lawrence L. Larmore. The server problem and on-line games. In *On-line algorithms (New Brunswick, NJ, 1991)*, volume 7 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 11–64. Amer. Math. Soc., Providence, RI, 1992.

[39] Marek Chrobak, Lawrence L. Larmore, Carsten Lund, and Nick Reingold. A better lower bound on the competitive ratio of the randomized 2-server problem. *Inf. Process. Lett.*, 63(2):79–83, 1997. doi:10.1016/S0020-0190(97)00099-9.

[40] Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006. doi:10.1002/047174882X.

[41] Imre Csiszár. The method of types. *IEEE Trans. Information Theory*, 44(6):2505–2523, 1998. doi:10.1109/18.720546.

[42] Erik D. Demaine, Dion Harmon, John Iacono, and Mihai Patrascu. Dynamic optimality - almost. *SIAM J. Comput.*, 37(1):240–251, 2007. doi:10.1137/S0097539705447347.

[43] Amir Dembo and Ofer Zeitouni. *Large deviations techniques and applications*, volume 38 of *Stochastic Modelling and Applied Probability*. Springer Berlin Heidelberg, corrected printing of the 1998 second edition, 2009. doi:http://dx.doi.org/10.1007/978-3-642-03311-7.

[44] Stefan Dobrev, Rastislav Královič, and Dana Pardubská. Measuring the problem-relevant information in input. *RAIRO - Theor. Inf. Appl.*, 43(3):585–613, 2009. doi:10.1051/ita/2009012.

[45] Christoph Dürr, Christian Konrad, and Marc P. Renault. On the power of advice and randomization for online bipartite matching. *CoRR*, abs/1602.07154, 2016.

[46] Sebastian Eggert, Lasse Kliemann, Peter Munstermann, and Anand Srivastav. Bipartite matching in the semi-streaming model. *Algorithmica*, 63(1-2):490–508, 2012. doi:10.1007/s00453-011-9556-8.

[47] Martin R. Ehmsen and Kim S. Larsen. Better bounds on online unit clustering. *Theor. Comput. Sci.*, 500:1–24, 2013. doi:10.1016/j.tcs.2013.07.008.

[48] Yuval Emek, Pierre Fraigniaud, Amos Korman, and Adi Rosén. Online computation with advice. *Theor. Comput. Sci.*, 412(24):2642–2656, 2011. doi:10.1016/j.tcs.2010.08.007.

[49] Leah Epstein and Rob van Stee. On the online unit clustering problem. *ACM Transactions on Algorithms*, 7(1):7, 2010. doi:10.1145/1868237.1868245.

[50] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004. doi:10.1016/j.jcss.2004.04.011.

[51] Xiequan Fan, Ion Grama, and Quansheng Liu. Hoeffding's inequality for supermartingales. *Stochastic Processes and their Applications*, 122(10):3545 – 3559, 2012.

[52] Amos Fiat, Dean P. Foster, Howard J. Karloff, Yuval Rabani, Yiftach Ravid, and Sundar Vishwanathan. Competitive algorithms for layered graph traversal. *SIAM J. Comput.*, 28(2):447–462, 1998. doi:10.1137/S0097539795279943.

[53] Amos Fiat, Richard M. Karp, Michael Luby, Lyle A. McGeoch, Daniel Dominic Sleator, and Neal E. Young. Competitive paging algorithms. *J. Algorithms*, 12(4):685–699, 1991. doi:10.1016/0196-6774(91)90041-V.

[54] Michal Forišek, Lucia Keller, and Monika Steinová. Advice complexity of online graph coloring. Unpublished manuscript, 2012.

[55] Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, 1997. doi:10.1006/jcss.1997.1504.

[56] Bernhard Fuchs, Winfried Hochstättler, and Walter Kern. Online matching on a line. *Theor. Comput. Sci.*, 332(1-3):251–264, 2005.

[57] Elisabeth Günther, Olaf Maurer, Nicole Megow, and Andreas Wiese. A new approach to online scheduling: Approximating the optimal competitive ratio. In *SODA*, pages 118–128. SIAM, 2013. doi:10.1137/1.9781611973105.9.

[58] Anupam Gupta and Kevin Lewi. The online metric matching problem for doubling metrics. In *ICALP (1)*, volume 7391, pages 424–435. Springer, 2012. doi:10.1007/978-3-642-31594-7_36.

[59] Sushmita Gupta, Shahin Kamali, and Alejandro López-Ortiz. On advice complexity of the k-server problem under sparse metrics. In *SIROCCO*, pages 55–67, 2013. doi:10.1007/978-3-319-03578-9_5.

[60] András Gyárfás and Jenö Lehel. On-line and first fit colorings of graphs. *Journal of Graph Theory*, 12(2):217–227, 1988. doi:10.1002/jgt.3190120212.

[61] Magnús M. Halldórsson and Mario Szegedy. Lower bounds for on-line graph coloring. *Theor. Comput. Sci.*, 130(1):163–174, 1994. doi:10.1016/0304-3975(94)90157-0.

[62] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

[63] Juraj Hromkovič, Rastislav Královič, and Richard Královič. Information complexity of online problems. In *MFCS*, pages 24–36, 2010. doi:10.1007/978-3-642-15155-2_3.

[64] Sandy Irani, Sandeep K. Shukla, and Rajesh Gupta. Algorithms for power savings. *ACM Transactions on Algorithms*, 3(4), 2007. doi:10.1145/1290672.1290678.

[65] Bala Kalyanasundaram and Kirk Pruhs. Online weighted matching. *J. Algorithms*, 14(3):478–488, 1993. doi:10.1006/jagm.1993.1026.

[66] Bala Kalyanasundaram and Kirk Pruhs. On-line network optimization problems. In Amos Fiat and Gerhard J. Woeginger, editors, *Online Algorithms, The State of the Art*, volume 1442 of *LNCS*, pages 268–280. Springer, 1996. doi:10.1007/BFb0029573.

[67] Anna R. Karlin, Mark S. Manasse, Lyle A. McGeoch, and Susan S. Owicki. Competitive randomized algorithms for non-uniform problems. In *SODA*, pages 301–309. SIAM, 1990. URL: http://dl.acm.org/citation.cfm?id=320176.320216.

[68] Anna R. Karlin, Mark S. Manasse, Larry Rudolph, and Daniel D. Sleator. Competitive snoopy caching. *Algorithmica*, 3:77–119, 1988. doi:10.1007/BF01762111.

[69] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358. ACM, 1990. doi:10.1145/100216.100262.

[70] Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theor. Comput. Sci.*, 127(2):255–267, 1994. doi:10.1016/0304-3975(94)90042-6.

[71] Hal A. Kierstead. Coloring graphs on-line. In *Online Algorithms, The State of the Art*, pages 281–305, 1996. doi:10.1007/BFb0029574.

[72] Dennis Komm, Rastislav Královič, Richard Královič, and Tobias Mömke. Randomized online algorithms with high probability guarantees. In *STACS*, volume 25 of *LIPIcs*, pages 470–481. Schloss Dagstuhl, 2014. doi:10.4230/LIPIcs.STACS.2014.470.

[73] Dennis Komm, Richard Královič, and Tobias Mömke. On the advice complexity of the set cover problem. In *CSR*, pages 241–252, 2012. doi:10.1007/978-3-642-30642-6_23.

[74] Elias Koutsoupias. The k-server problem. *Computer Science Review*, 3(2):105–118, 2009. doi:10.1016/j.cosrev.2009.04.002.

[75] Elias Koutsoupias and Christos H. Papadimitriou. On the k-server conjecture. *J. ACM*, 42(5):971–983, 1995.

[76] H. W. Kuhn. Extensive games and the problem of information. In *Contributions to the theory of games, vol. 2*, Annals of Mathematics Studies, no. 28, pages 193–216. Princeton University Press, 1953.

[77] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Inf. Comput.*, 108(2):212–261, 1994. Preliminary version in FOCS'89. doi:10.1006/inco.1994.1009.

[78] Carsten Lund and Nick Reingold. Linear programs for randomized on-line algorithms. In Daniel Dominic Sleator, editor, *SODA*, pages 382–391. ACM/SIAM, 1994.

[79] Mark S. Manasse, Lyle A. McGeoch, and Daniel Dominic Sleator. Competitive algorithms for on-line problems. In *STOC*, pages 322–333. ACM, 1988. doi:10.1145/62212.62243.

[80] M. Maschler, S. Zamir, and E. Solan. *Game Theory*. Cambridge University Press, 2013.

[81] Adam Meyerson, Akash Nanavati, and Laura J. Poplawski. Randomized online algorithms for minimum metric bipartite matching. In *SODA*, pages 954–959. ACM Press, 2006. URL: http://dl.acm.org/citation.cfm?id=1109557.1109662.

[82] Jesper W. Mikkelsen. Optimal online edge coloring of planar graphs with advice. In *CIAC*, volume 9079 of *Lecture Notes in Computer Science*, pages 352–364. Springer, 2015. doi:10.1007/978-3-319-18173-8_26.

[83] Shuichi Miyazaki. On the advice complexity of online bipartite matching and online stable marriage. *Inf. Process. Lett.*, 114(12):714–717, 2014.

[84] Tobias Mömke. A competitive ratio approximation scheme for the k-server problem in fixed finite metrics. *CoRR*, abs/1303.2963, 2013. URL: http://arxiv.org/abs/1303.2963.

[85] Ilan Newman. Private vs. common random bits in communication complexity. *Inf. Process. Lett.*, 39(2):67–71, 1991. doi:10.1016/0020-0190(91)90157-D.

[86] Harald Räcke, Christian Sohler, and Matthias Westermann. Online scheduling for sorting buffers. In *ESA*, volume 2461 of *Lecture Notes in Computer Science*, pages 820–832. Springer, 2002. doi:10.1007/3-540-45749-6_71.

[87] Marc P. Renault and Adi Rosén. On online algorithms with advice for the k-server problem. *Theory Comput. Syst.*, 56(1):3–21, 2015. doi:10.1007/s00224-012-9434-z.

[88] Daniel D. Sleator and Robert E. Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, 1985. doi:10.1145/2786.2793.

[89] Daniel Dominic Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *J. ACM*, 32(3):652–686, 1985. doi:10.1145/3828.3835.

[90] Jasmin Smula. *Information Content of Online Problems, Advice versus Determinism and Randomization*. PhD thesis, ETH Zürich, 2015. doi:10.3929/ethz-a-010497710.

[91] Boris Teia. A lower bound for randomized list update algorithms. *Inf. Process. Lett.*, 47(1):5–9, 1993. doi:10.1016/0020-0190(93)90150-8.

[92] Robert E. Wilber. Lower bounds for accessing binary search trees with rotations. *SIAM J. Comput.*, 18(1):56–67, 1989. doi:10.1137/0218004.

[93] Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *FOCS*, pages 222–227, 1977. doi:10.1109/SFCS.1977.24.

A Derandomization using advice

In this appendix, we slightly extend the derandomization result due to Böckenhauer et al. [25] (see also the survey [23]). In particular, we show that it is possible to derandomize algorithms which uses both advice and randomization. Also, we show how to modify the result for minimization problems so that it can be applied to maximization problems. Finally, we provide lower bounds showing that these derandomization results are essentially best possible.

Theorem 16. *Let* P *be an online minimization problem and let* R *be a randomized algorithm with advice complexity* $b(n)$*. Suppose that there exists a function* $I : \mathbb{N} \to \mathbb{N}$ *such that the number of valid* P*-inputs of length* n *is at most* $I(n)$*. Then for every* $\varepsilon > 0$*, there exists a deterministic algorithm with advice,* `ALG`*, such that* $\texttt{ALG}(\sigma) \leq (1+\varepsilon)\,\mathbb{E}[R(\sigma)]$ *for every input* σ *and such that* `ALG` *reads at most*

$$b(n) + O(\log n + \log\log I(n) + \log \varepsilon^{-1})$$

bits of advice on inputs of length n*. In particular, if* $I(n) = 2^{n^{O(1)}}$ *and* $\varepsilon > 0$ *is a small but fixed constant, then* `ALG` *reads* $b(n) + O(\log n)$ *bits of advice.*

Proof. Let I_n be the set of all P-inputs of length n. By assumption, $|I_n| \leq I(n)$. Recall that a randomized algorithm with advice complexity b is a probability distribution over deterministic algorithms with advice complexity (at most) b. Let t be an integer and let the random variables $R_1, \ldots, R_t$ be t deterministic algorithms selected independently at random according to the probability distribution specified by the randomized algorithm R. Then, for every fixed input $\sigma \in I_n$ such that $\mathbb{E}[R(\sigma)] > 0$ and every $1 \leq i \leq t$, Markov's inequality implies that

$$\Pr\Big[R_i(\sigma) > (1+\varepsilon)\,\mathbb{E}[R(\sigma)]\Big] \leq \frac{1}{1+\varepsilon}. \tag{56}$$

Recall that costs are by definition non-negative (Definition 1). Therefore, if $\mathbb{E}[R(\sigma)] = 0$ for some $\sigma \in I_n$, then $R_i(\sigma) = 0$ for $1 \leq i \leq t$. Thus, for every $\sigma \in I_n$ where $\mathbb{E}[R(\sigma)] = 0$, the inequality (56) holds for the trivial reason that the probability on the left-hand side is zero.

Since $R_1, \ldots, R_t$ are independent random variables, inequality (56) implies that for every fixed input $\sigma \in I_n$, the probability of all t algorithms failing to achieve the desired guarantee is at most

$$\Pr\Big[\min_{1\leq i\leq t}\{R_i(\sigma)\} > (1+\varepsilon)\,\mathbb{E}[R(\sigma)]\Big] \leq \left(\frac{1}{1+\varepsilon}\right)^t.$$

An application of the union bound gives

$$\Pr\Big[\exists \sigma \in I_n : \min_{1\leq i\leq t}\{R_i(\sigma)\} > (1+\varepsilon)\,\mathbb{E}[R(\sigma)]\Big] \leq I(n)\left(\frac{1}{1+\varepsilon}\right)^t. \tag{57}$$

Note that

$$I(n)\left(\frac{1}{1+\varepsilon}\right)^t < 1 \Leftrightarrow \frac{\log(I(n))}{\log(1+\varepsilon)} < t.$$

Choose $t = t_{n,\varepsilon}$ to be the smallest integer such that $\frac{\log(I(n))}{\log(1+\varepsilon)} < t_{n,\varepsilon}$. Then

$$\Pr\Big[\forall \sigma \in I_n : \min_{1\leq i\leq t_{n,\varepsilon}}\{R_i(\sigma)\} \leq (1+\varepsilon)\,\mathbb{E}[R(\sigma)]\Big] \geq 1 - I(n)\left(\frac{1}{1+\varepsilon}\right)^{t_{n,\varepsilon}} > 0.$$

This shows that there exists $t_{n,\varepsilon}$ deterministic algorithms such that for every input $\sigma \in I_n$, at least one of these $t_{n,\varepsilon}$ algorithms incurs a cost of at most $(1+\varepsilon)\,\mathbb{E}[R(\sigma)]$. By assumption, the advice complexity of each of these $t_{n,\varepsilon}$ algorithms is at most $b(n)$.

We are now ready to define $\texttt{ALG}$. The oracle will write the length, n, of the input onto the advice tape using $O(\log n)$ bits. Knowing n, $\texttt{ALG}$ will be able to[11] determine $t_{n,\varepsilon}$ algorithms $\texttt{ALG}_1, \ldots, \texttt{ALG}_{t_{n,\varepsilon}}$ such that for every input $\sigma \in I_n$, it holds that $\min_i\{\texttt{ALG}_i(\sigma)\} \leq (1+\varepsilon)\,\mathbb{E}[R(\sigma)]$. The oracle will encode the index, i, of the best of these $t_{n,\varepsilon}$ algorithms for the current input σ. This requires $\lceil \log t_{n,\varepsilon} \rceil$ bits. Furthermore, the oracle writes the advice read by $\texttt{ALG}_i$ (when given σ as input) onto the advice tape of $\texttt{ALG}$. The algorithm $\texttt{ALG}$ is then able to simulate $\texttt{ALG}_i$ and, thus, $\texttt{ALG}(\sigma) = \texttt{ALG}_i(\sigma) \leq (1+\varepsilon)\,\mathbb{E}[R(\sigma)]$. The number of advice bits read by $\texttt{ALG}$ is at most

$$b(n) + \left\lceil \log\left(\frac{\log(I(n))}{\log(1+\varepsilon)} + 1\right)\right\rceil + O(\log n) = b(n) + O(\log n + \log\log I(n) + \log \varepsilon^{-1}).$$

□

Note that ε could be non-constant. Assuming $I(n) = 2^{n^{O(1)}}$, we can take, for example, $\varepsilon = 1/n$ and get an algorithm $\texttt{ALG}$ which reads $O(\log n)$ bits of advice while $\texttt{ALG}(\sigma) \leq (1+1/n)\,\mathbb{E}[R(\sigma)]$.

We will now show how to obtain a guarantee on the competitive ratio using Theorem 16.

Corollary 1. *Let $c \geq 1$ be a constant, let P be a minimization problem and let R be a randomized c-competitive algorithm with advice complexity $b(n)$. Suppose that there exists a function $I : \mathbb{N} \to \mathbb{N}$ such that the number of valid P-inputs of length n is at most $I(n)$. If $I(n) = 2^{n^{O(1)}}$ then for every constant $\varepsilon > 0$, there exists a deterministic $(c+\varepsilon)$-competitive algorithm with advice complexity $b(n) + O(\log n)$.*

Proof. Choose $\varepsilon' = \varepsilon/c$ and apply Theorem 16 to R. This gives a deterministic algorithm $\texttt{ALG}$ with the desired advice complexity such that $\texttt{ALG}(\sigma) \leq (1+\varepsilon')\,\mathbb{E}[R(\sigma)]$ for every input σ. Using that R is c-competitive, we get that for every input σ,

$$\texttt{ALG}(\sigma) \leq (1+\varepsilon')\,\mathbb{E}[R(\sigma)] \leq (1+\varepsilon')(c \cdot \texttt{OPT}(\sigma) + \alpha) \leq (c+\varepsilon)\texttt{OPT}(\sigma) + (1+\varepsilon)\alpha.$$

This proves that $\texttt{ALG}$ is $(c+\varepsilon)$-competitive. □

We will now prove a derandomization result for maximization algorithms. However, there are some technical difficulties involved in this. For a minimization problem, we always know that the cost of a solution is at least zero. For a maximization problem, $\texttt{OPT}(\sigma)$ could grow arbitrarily fast (as a function of the length of the input) while the profit of a randomized algorithm could remain small. We resolve this by assuming that the strict competitive ratio of the algorithm is bounded by some function $c(n)$, and show that if $c(n)$ is polynomial in n, then we get the same conclusion as for minimization problems. In Theorem 19, we show that this assumption cannot be removed (without replacing it with something else).

[11] Exactly how this is done depends on the model of online algorithms being used. If an online algorithm is just a strategy for a request-answer game without computational constraints (such as in the model of Ben-David et al. [17], see also Appendix C.1), then we do not need to specify how $\texttt{ALG}$ determines these algorithms (it suffice to know that they exist). In most other models, such as the one used in [25], it will be possible for $\texttt{ALG}$ to use exhaustive search.

Theorem 17. *Let* P *be an online maximization problem and let* R *be a randomized algorithm with advice complexity* $b(n)$*. Suppose that there exists a function* $I : \mathbb{N} \to \mathbb{N}$ *such that the number of valid* P*-inputs of length* n *is at most* $I(n)$*. Furthermore, suppose that* R *is strictly* $c(n)$*-competitive, that is,* $\texttt{OPT}(\sigma) \leq c(n) \cdot E[R(\sigma)]$ *for every input* σ *of length at most* n*.*

Then for every $0 < \varepsilon < 1$*, there exists a deterministic algorithm with advice,* `ALG`*, such that* $\texttt{ALG}(\sigma) \geq (1-\varepsilon)\,\mathbb{E}[R(\sigma)]$ *for every input* σ *and such that* `ALG` *reads at most*

$$b(n) + O(\log n + \log\log I(n) + \log c(n) + \log \varepsilon^{-1})$$

bits of advice on inputs of length n*. In particular, if* $I(n) = 2^{n^{O(1)}}$*,* $c(n) = n^{O(1)}$*, and* ε *is a small but fixed constant, then* `ALG` *reads* $b(n) + O(\log n)$ *bits of advice.*

Proof. Consider the set I_n of all P-inputs of length n. By assumption, $|I_n| \leq I(n)$. Recall that a randomized algorithm with advice complexity b is a probability distribution over deterministic algorithms with advice complexity (at most) b. Let t be an integer and let the random variables $R_1, \ldots, R_t$ be t deterministic algorithms selected independently at random according to the probability distribution specified by the randomized algorithm R. Then, for every fixed input $\sigma \in I_n$ such that $\mathbb{E}[R(\sigma)] > 0$ and every $1 \leq i \leq t$, Markov's inequality (applied to the positive random variable $\texttt{OPT}(\sigma) - R_i(\sigma)$) implies that

$$\begin{aligned}
\Pr\Big[R_i(\sigma) < (1-\varepsilon)\,\mathbb{E}[R(\sigma)]\Big] &\leq \frac{\texttt{OPT}(\sigma) - \mathbb{E}[R(\sigma)]}{\texttt{OPT}(\sigma) - (1-\varepsilon)\,\mathbb{E}[R(\sigma)]} \\
&= \frac{\texttt{OPT}(\sigma)/\,\mathbb{E}[R(\sigma)] - 1}{\texttt{OPT}(\sigma)/\,\mathbb{E}[R(\sigma)] - (1-\varepsilon)} \\
&\leq \frac{c(n) - 1}{c(n) - (1-\varepsilon)}. \qquad (58)
\end{aligned}$$

The same upper bound trivially holds if $E[R(\sigma)] = 0$, since then the probability on the left-hand side of (58) is zero (because profits are by definition non-negative). Since $R_1, \ldots, R_t$ are independent random variables, inequality (58) implies that for every fixed input $\sigma \in I_n$, the probability of all t algorithms failing to achieve the desired performance guarantee is at most

$$\Pr\Big[\min_{1\leq i\leq t}\{R_i(\sigma)\} < (1-\varepsilon)\,\mathbb{E}[R(\sigma)]\Big] \leq \left(\frac{c(n)-1}{c(n)-(1-\varepsilon)}\right)^t.$$

An application of the union bound gives

$$\Pr\Big[\exists \sigma \in I_n : \min_{1\leq i\leq t}\{R_i(\sigma)\} < (1-\varepsilon)\,\mathbb{E}[R(\sigma)]\Big] \leq I(n)\left(\frac{c(n)-1}{c(n)-(1-\varepsilon)}\right)^t. \qquad (59)$$

Note that

$$I(n)\cdot\left(\frac{c(n)-1}{c(n)-(1-\varepsilon)}\right)^t < 1 \Leftrightarrow \frac{\log(I(n))}{\log\left(\frac{c(n)-(1-\varepsilon)}{c(n)-1}\right)} < t.$$

Using that $\log(1+x)^{-1} \leq 1/x$ for $0 < x \leq 1$ and $\log(1+x)^{-1} \leq 1$ for $x > 1$, we get that

$$\begin{aligned}
\log\left(\frac{c(n)-(1-\varepsilon)}{c(n)-1}\right)^{-1} &= \log\left(1 + \frac{\varepsilon}{c(n)-1}\right)^{-1} \\
&\leq \frac{c(n)-1}{\varepsilon} + 1.
\end{aligned}$$

Choose $t = t_{n,\varepsilon}$ to be the smallest integer such that $\log I(n) \cdot \left(\frac{c(n)-1}{\varepsilon} + 1\right) < t_{n,\varepsilon}$. Then

$$\Pr\left[\forall \sigma \in I_n : \min_{1 \le i \le t_{n,\varepsilon}} \{R_i(\sigma)\} \ge (1-\varepsilon)\,\mathbb{E}[R(\sigma)]\right] > 0.$$

This shows that there exist $t_{n,\varepsilon}$ deterministic algorithms such that for every input $\sigma \in I_n$, at least one of these $t_{n,\varepsilon}$ algorithms obtains a profit of at least $(1-\varepsilon)\,\mathbb{E}[R(\sigma)]$. By assumption, the advice complexity of each of these $t_{n,\varepsilon}$ algorithms is at most $b(n)$.

We are now ready to define ALG. The oracle will write the length n of the input onto the advice tape using $O(\log n)$ bits. Knowing n, ALG will be able to determine $t_{n,\varepsilon}$ algorithms $\mathtt{ALG}_1, \ldots, \mathtt{ALG}_t$ such that $\min_i\{\mathtt{ALG}_i(\sigma)\} \ge (1-\varepsilon)\,\mathbb{E}[R(\sigma)]$ for every $\sigma \in I_n$. The oracle will encode the index, i, of the best of the $t_{n,\varepsilon}$ algorithms for the current input σ. This requires $\lceil \log t_{n,\varepsilon} \rceil$ bits. Furthermore, the oracle writes the advice read by $\mathtt{ALG}_i$ (when given σ as input) onto the advice tape of ALG. The algorithm ALG is then able to simulate $\mathtt{ALG}_i$ and, thus, $\mathtt{ALG}(\sigma) = \mathtt{ALG}_i(\sigma) \ge (1-\varepsilon)\,\mathbb{E}[R(\sigma)]$. The number of advice bits read by ALG is at most

$$\begin{aligned} &b(n) + \log\left(\log I(n) \cdot \left(\frac{c(n)-1}{\varepsilon} + 1\right) + 1\right) + O(\log n) \\ &\quad = b(n) + O(\log n + \log\log I(n) + \log c(n) + \log \varepsilon^{-1}). \end{aligned}$$

□

Corollary 2. *Let $c \ge 1$ be a constant. Furthermore, let* P *be an online maximization problem and let R be a randomized algorithm with advice complexity $b(n)$. Suppose that there exists a function $I : \mathbb{N} \to \mathbb{N}$ such that the number of valid* P*-inputs of length n is at most $I(n)$. If $I(n) = 2^{n^{O(1)}}$, then for every $\varepsilon > 0$, if there exists a c-competitive randomized algorithm R with advice complexity $b(n)$, then there exists a $(c+\varepsilon)$-competitive deterministic algorithm with advice complexity $b(n)+O(\log n)$.*

Proof. By assumption, there exists an additive constant α such that $\mathtt{OPT}(\sigma) \le c\,\mathbb{E}[R(\sigma)]+\alpha$ for every input σ. Let $I_{2\alpha}$ those P-inputs for which $2\alpha \le \mathtt{OPT}(\sigma)$. Note that if $\sigma \in I_{2\alpha}$, then $c\,\mathbb{E}[R(\sigma)] + \alpha \ge \mathtt{OPT}(\sigma) \ge 2\alpha$. Thus, $c\,\mathbb{E}[R(\sigma)] \ge \alpha$, and hence $2c\,\mathbb{E}[R(\sigma)] \ge \mathtt{OPT}(\sigma)$ for every $\sigma \in I_{2\alpha}$. In particular, R is strictly $2c$-competitive when we restrict to inputs from $I_{2\alpha}$. Choose $\varepsilon' = \varepsilon/(c+\varepsilon)$. By applying Theorem 17 to the algorithm R restricted to inputs from $I_{2\alpha}$, we obtain a deterministic algorithm $\mathtt{ALG}'$ with advice complexity $b(n) + O(\log n)$ such that $\mathtt{ALG}'(\sigma) \ge (1-\varepsilon)\,\mathbb{E}[R(\sigma)]$ for every $\sigma \in I_{2\alpha}$.

Our deterministic $(c+\varepsilon)$-competitive algorithm ALG works as follows: A single bit of advice is used to indicate if $\sigma \in I_{2\alpha}$ or not. If $\sigma \notin I_{2\alpha}$, then ALG computes an arbitrary valid output. If $\sigma \in I_{2\alpha}$, the algorithm simulates $\mathtt{ALG}'$ on σ. By the choice of ε', we get that for every $\sigma \in I_{2\alpha}$,

$$\mathtt{OPT}(\sigma) \le c\,\mathbb{E}[R(\sigma)] + \alpha \le \frac{c}{1-\varepsilon'}\mathtt{ALG}'(\sigma) + \alpha = (c+\varepsilon)\mathtt{ALG}'(\sigma) + \alpha = (c+\varepsilon)\mathtt{ALG}(\sigma) + \alpha.$$

Since $\mathtt{OPT}(\sigma) \le 2\alpha$ for every $\sigma \notin I_{2\alpha}$, it follows that ALG is $(c+\varepsilon)$-competitive with an additive constant of 2α. □

A.1 Lower bounds for derandomization using advice

In this section, we show that Theorem 16 is essentially tight by providing (for any choice of $I(n)$) an online problem where $\Omega(\log\log I(n))$ bits of advice are necessary for a deterministic online algorithm to achieve a competitive ratio anywhere near that of the best randomized algorithm.

Theorem 18. *Let $I(n) : \mathbb{N} \to \mathbb{N}$ be an increasing function. There exists a minimization problem* P *with at most $I(n)$ inputs of length n such that a randomized algorithm without advice can be* 1*-competitive, while a deterministic algorithm with advice needs to read at least $\Omega(\log \log I(n))$ bits of advice to achieve a constant competitive ratio.*

Proof. Define $f(n) = \lfloor \sqrt{\log I(n)} \rfloor$ and $u(n) = 2^{f(n)}$. Furthermore, for each $n \geq 2$, let M_n be a set containing $u(n)$ elements. The problem P is defined as follows: In the first $n-2$ rounds, a dummy request appears which requires no response. In round $n-1$, a special request, ?, is revealed and the algorithm must respond by choosing an element x from M_n (note that the algorithm can deduce n from the number of dummy requests). In round n, a set $M' \subset M_n$ such that $|M'| = f(n)$ is revealed (no response is required from the algorithm). If $x \in M'$, the algorithm incurs a cost of $f(n)$. Otherwise, the algorithm incurs a cost of 1. Formally, the input to P is a sequence $(d_1, \ldots, d_{n-2}, ?, M')$ where d_i denotes a dummy request. Note that since $f(n) < u(n)$, it follows that $\texttt{OPT}(\sigma) = 1$ for any input. Also, the number of inputs of length n is at most $\binom{u(n)}{f(n)} \leq u(n)^{f(n)} = 2^{f(n)^2} \leq 2^{\log I(n)} = I(n)$.

Consider the randomized algorithm R which chooses x uniformly at random from M_n in round $n-1$. Since the adversary is oblivious, the expected cost incurred by R on an input σ of length n is

$$\mathbb{E}[R(\sigma)] \leq f(n) \cdot \frac{f(n)}{u(n)} + 1 \cdot \left(1 - \frac{f(n)}{u(n)}\right) \leq \frac{f(n)^2}{u(n)} + 1 \leq 3.$$

It follows that $\mathbb{E}[R(\sigma)] \leq \texttt{OPT}(\sigma) + 2$ and hence R is 1-competitive.

Let $\texttt{ALG}$ be a deterministic algorithm reading at most $b(n)$ bits of advice on inputs of length n. Since the first $n-1$ request are identical in all inputs, the element x chosen in round $n-1$ can only depend on n and the advice. Thus, $\texttt{ALG}$ can choose at most $2^{b(n)}$ different elements from M_n on inputs of length n. If for some n we have that $2^{b(n)} \leq f(n)$, then an adversary can choose M' to contain all of the elements that could possibly be outputted by $\texttt{ALG}$. Thus, $\texttt{ALG}$ will incur a cost of $f(n)$. In particular, if $\texttt{ALG}$ is to achieve a constant competitive ratio, then there must be an N such that for all $n \geq N$, we have that $2^{b(n)} > f(n)$ and, hence,

$$b(n) > \log(f(n)) = \log\left(\lfloor \sqrt{\log I(n)} \rfloor\right) = \Omega(\log \log I(n)).$$

□

For maximization problems, it is natural to ask if the $O(\log c(n))$-term in Theorem 17 is necessary or just an artifact of our proof-strategy. We show in the following example that it is indeed necessary, by designing an online problem and a strictly c-competitive randomized algorithm such that any deterministic algorithm reading $o(\log c(n))$ bits of advice will, on some inputs, obtain a profit much smaller than the randomized algorithm.

Theorem 19. *Let $c : \mathbb{N} \to \mathbb{N}$ be a function. There exists an online maximization problem* P *with $I(n)$ inputs of length n and a randomized* P*-algorithm R without advice such that:*

(i) The algorithm R is strictly $c(n)$-competitive and $\mathbb{E}[R(\sigma)] = n$ for every input σ of length n.

(ii) If $\texttt{ALG}$ is a deterministic algorithm reading $O(\log n + \log \log I(n))$ bits of advice, then there exists an n_0 such that for every $n \geq n_0$, there is an input σ of length n such that $\texttt{ALG}(\sigma) = 0$.

(iii) If $\texttt{ALG}$ is a deterministic algorithm such that $\texttt{ALG}(\sigma) \geq 1$ for every input σ, then $\texttt{ALG}$ must read at least $\Omega(\log c(n))$ bits of advice.

Proof. Let P be the following modification of the string guessing problem with unknown history (cf. [24]): The problem P consists of n rounds. The number of rounds is revealed to the online algorithm at the beginning. In round $1 \leq i \leq n$, the algorithm has to guess a character from the alphabet $\{1, 2, \ldots, c(n)\}$. Only after all n rounds, the correct characters are revealed (thus, the algorithm does not learn the correct character right after guessing). For each character correctly guessed by the algorithm, it obtains a profit of $c(n)$. For each wrong guess, the algorithm obtains a profit of 0.

Note that $\texttt{OPT}(\sigma) = n \cdot c(n)$ and that $I(n) = c(n)^n$. Let R be the algorithm which selects a character uniformly at random in each round. The probability of R guessing right in round i is $1/c(n)$. Since the profit of a correct guess is $c(n)$, this means that the expected profit of R in each round is 1. By linearity of expectation, $\mathbb{E}[R(\sigma)] = n$ for every input σ of length n. It follows that R is strictly $c(n)$-competitive.

Let $\texttt{ALG}$ be a deterministic algorithm with advice complexity $b(n)$. Suppose that for some n, it holds that $b(n) \leq \log c(n) - 1$. This means that $\texttt{ALG}$ can produce at most $2^{b(n)} \leq 2^{\log(c(n))-1} = c(n)/2 < c(n)$ different outputs on inputs of length n. In particular, for each $1 \leq i \leq n$, there exists a character $x_i \in \{1, 2, \ldots, c(n)\}$ such that none of the outputs that $\texttt{ALG}$ can produce has x_i as the answer in round i. It follows that there exists an input string $x = x_1 \ldots x_n$ of length n such that $\texttt{ALG}(x) = 0$. This proves that (iii) holds. In order to see that (ii) holds, note that $\log \log I(n) = \log \log(c(n)^n) = O(\log n + \log \log c(n))$. In particular, if $b(n) = O(\log n + \log \log I(n))$, then $b(n) = o(\log c(n))$ which shows that (ii) holds.

□

B Concentration inequalities

B.1 Chernoff bounds

Theorem 20 ([41, Lemma II.2])**.** *Let $0 < p < 1$ and let $X_1, X_2, \ldots$ be a sequence of independent random variables such that for each i, $\Pr[X_i = 1] = p$ and $\Pr[X_i = 0] = 1 - p$. Then for every $\varepsilon > 0$,*

$$\frac{2^{-K_p(p+\varepsilon)\cdot n}}{n+1} \leq \Pr\left[\sum_{i=1}^{n} X_i \geq (p+\varepsilon)n\right] \leq 2^{-K_p(p+\varepsilon)\cdot n}. \tag{60}$$

The upper bound in (60) is a variant of the Chernoff bound. The lower bound in (60) shows that this variant of the Chernoff bound is essentially tight. (60) is well-known and can, for example, be easily derived as a special case of Lemma II.2 in [41].

B.2 The Azuma-Hoeffding inequality

The Azuma-Hoeffding inequality [11, 62] gives a bound similar to (60) for bounded-difference martingales.

Theorem 21 ([43, 51])**.** *Let $Z_0, Z_1, \ldots Z_n$ be a (real-valued) supermartingale. Assume that there exist constants $d, \sigma \geq 0$ such that, for $0 \leq i < n$,*

$$Z_{i+1} - Z_i \leq d,$$
$$\mathbb{E}[(Z_{i+1} - Z_i)^2 | Z_1, \ldots, Z_i] \leq \sigma^2.$$

Then, for every $r \geq 0$,

$$\Pr[Z_n - Z_0 \geq tn] \leq \exp_2\left(-K_{\frac{\gamma}{1+\gamma}}\left(\frac{\alpha+\gamma}{1+\gamma}\right)n\right), \tag{61}$$

where $\gamma = \dfrac{\sigma^2}{d^2}$ *and* $\alpha = \dfrac{t}{d}$ *(and* $\exp_2(x) = 2^x$*).*

(61) is almost identical to Corollary 2.4.7 in [43], except that the corollary in [43] is only stated for martingales, not supermartingales. However, one may verify that the same proof applies to supermartingales. Alternatively, (61) can be derived from Theorem 2.1 in [51], which is stated and proved for supermartingales. Indeed, under the assumptions of (61), it follows immediately from [51] that

$$\Pr[Z_n - Z_0 \geq tn] \leq \left(\left(\frac{\gamma}{\alpha+\gamma}\right)^{n(\alpha+\gamma)}\left(\frac{1}{1-\alpha}\right)^{n(1-\alpha)}\right)^{\frac{1}{1+\gamma}}. \tag{62}$$

where $\gamma = \sigma^2/d^2$ and $\alpha = t/d$. It can be shown by a straightforward algebraic calculation that the right-hand side of (62) is identical to the right-hand side of (61).

(61) provides a very good bound (for example, it can be used to obtain the upper bound of (60)) but it is sometimes a little cumbersome to work with. If one is willing to settle for a weaker bound, the following version of Azuma-Hoeffding is easier to use:

Theorem 22 ([11])**.** *Let* $Z_0, Z_1, \ldots Z_n$ *be a (real-valued) supermartingale. Assume that there exists a constant* d *such that for* $0 \leq i < n$, $|Z_{i+1} - Z_i| \leq d$*. Then*

$$\Pr[Z_n - Z_0 \geq tn] \leq \exp\left(-\frac{t^2}{2d^2}n\right). \tag{63}$$

C Models of online computation

In this appendix, we will formally define what we mean by a randomized online algorithm with advice.

C.1 Randomized online algorithms

At STOC'90, Ben-David, Borodin, Karp, Tardos, and Wigderson presented a formal model of deterministic and randomized online algorithms, and proved several fundamental results regarding this model [17]. We will refer to this model as the *BBKTW-model*. The BBKTW-model became a popular model for formalizing randomized online computation. For example, the model is used in the textbook by Borodin and El-Yaniv [27]. In the BBKTW-model, a deterministic online algorithm is simply a set of mappings which defines a strategy for a Request-Answer game (see [17] for the details). In particular, no computational restrictions are imposed. The model only considers what is mathematically possible to achieve in an online environment with incomplete information and does not take computational resources into account. A randomized online algorithm is then defined as a probability distribution over deterministic online algorithms.

A probability distribution over deterministic algorithms corresponds to a *mixed strategy*. Another very natural way to define a randomized algorithm is to allow the online algorithm to flip a

number of coins before answering each request. Such a definition of a randomized online algorithms corresponds to a *behavioral strategy*. Note that in an online problem, the algorithm has complete information of all previous requests and answers. It follows from Kuhn's theorem [76, 7] that because of this complete information, mixed and behavioral strategies are equivalent. Intuitively, one can think of the probability distribution of a mixed strategy as fixing the coin-flips of the algorithm. In the case of finite online problems, the textbook [27] by Borodin and El-Yaniv contains a full proof of this equivalence. However, online problems are usually not finite, since the length of the input is unbounded. For an accessible proof of the equivalence between mixed and behavioral strategies for infinite games, we refer to e.g. [80].

One possible alternative to our definition of a randomized algorithm with advice would be to define the advice complexity as the *expected* number of advice bits read by the algorithm over inputs of length at most n. We suspect that most of our lower bounds would also hold for such slightly more powerful randomized algorithms with advice but leaves this as an open question.

C.2 The advice-at-beginning model

Recall that according to Definition 1, the advice is provided to the online algorithm on an infinite tape in order to prevent the algorithm to learn anything from the length of the advice. We will now introduce another model of online algorithms with advice in which the advice is given as a finite bitstring. All of our lower bounds holds in this slightly more powerful model.

Definition 14 (Advice-at-beginning)**.** An online algorithm with advice in the *advice-at-beginning* model is defined as a triple $\texttt{ALG} = \left(\{D_x\}_{x\in\{0,1\}^*}, b, \varphi\right)$ where $b : \mathbb{N} \to \mathbb{N} \cup \{0\}$ is a non-decreasing function, $\varphi : I \to \{0,1\}^*$ is a mapping such that $|\varphi(\sigma)| = b(n)$ where $n = |\sigma|$ is the number of requests in σ, and, for each bit-string $x \in \{0,1\}^*$, D_x is a deterministic online algorithm in the Request-Answer model. On input σ, the algorithm $\texttt{ALG}$ produces the same output as the deterministic algorithm $D_{\varphi(\sigma)}$. In particular, $\text{score}(\texttt{ALG}(\sigma)) = \text{score}(D_{\varphi(\sigma)}(\sigma))$. The function $b(n)$ is called the *advice complexity* of $\texttt{ALG}$.

Please note that we do not allow the algorithm to receive as advice a bitstring which is shorter than $b(n)$. This is in order to ensure that an algorithm with advice complexity $b(n)$ corresponds to at most $2^{b(n)}$ different deterministic algorithms. If the advice string was allowed to be smaller than $b(n)$, the algorithm would correspond to $\sum_{i=0}^{b(n)} 2^i = 2^{b(n)+1} - 1$ different deterministic algorithms. This is not particularly important but slightly simplifies the model.

Relationship to the advice-on-tape model. An algorithm in the advice-on-tape model using b bits of advice can be converted to an algorithm in the advice-at-beginning model using b bits of advice. This is why a lower bound in the advice-at-beginning model immediately translates into a lower bound in the advice-on-tape model.

An algorithm in the advice-at-beginning model using b bits of advice can be converted to an algorithm in the advice-on-tape model as follows: On input σ, the oracle computes $b(n)$ where $n = |\sigma|$ and it computes $\varphi(\sigma)$. It then writes the integer $b(n)$ onto the advice tape in a self-delimiting way followed by the string $\varphi(\sigma)$. This requires at most $b + 2\lceil\log(b+1)\rceil + 1 = b + O(\log b)$ bits of advice. This shows that the two models are equivalent up to an additive logarithmic difference in the advice complexity.

D Modeling online problems as task systems

In this appendix, we show that the following online problems can be modeled as a finite lazy task system (Definition 10):

1. The MTS problem on a finite metric space.
2. The k-server problem on a finite metric space.
3. The list update problem.
4. The dynamic binary search three problem on a set of N nodes.

We also show that reordering buffer management is Σ-repeatable.

MTS, k-server, and list update. If d is a metric, then d also satisfies the conditions imposed on the distance function in the definition of a lazy task system. This proves the following lemma.

Lemma 12. *If $(\mathcal{S}, \mathcal{T})$ is a MTS on a finite metric space, then $(\mathcal{S}, \mathcal{T})$ is also a finite lazy task system.*

It is easy to see that the k-server problem can be modeled as a MTS. It is shown in e.g. [27] how to model the list update problem as a MTS.

Lemma 13. *The k-server problem on a finite metric space can be modeled as a finite lazy task system.*

Lemma 14. *The list update problem can be modeled as a finite lazy task system.*

Dynamic binary search trees. We will show how to model the dynamic binary search tree problem as a lazy task system (when defined as in Section 6.7).

Lemma 15. *The dynamic binary search three problem on a set of N nodes can be modeled as a finite lazy task system.*

Proof. Fix N nodes. The set of states $\mathcal{S}$ consists of all possible states s of the form $s = (T, v_r, v_e)$ where T is a binary search tree on N nodes and v_r, v_e are nodes in T. We say that v_r is the *request node* of s and that v_e is the *end note* of s. We will now define the set of tasks $\mathcal{T}$. For each node v, there is an associated task t_v defined as follows:

$$t_v(s) = \begin{cases} 0 & \text{if } v \text{ is the request node of } s \\ \infty & \text{if } v \text{ is not the request node of } s \end{cases}$$

The distance $d(s, s')$ between two states $s = (T, v_r, v_e), s' = (T', v'_r, v'_e)$ is defined as follows: $d(s, s')$ is the minimum number of unit-cost operations that must be performed such that the following conditions are satisfied:

1. The pointer is initialized at the node v_e in the tree T (recall that initialization is a unit-cost operation).
2. The rotations performed during the sequence of unit-cost operations transforms T into T'.
3. The node v'_r is visited during the sequence of unit-cost operations.

4. The pointer ends at the node v'_e in the tree T'.

We will argue that the LTS defined above models the binary search tree problem. Consider an input sequence $\sigma = (x_1, \ldots, x_n)$ and an initial state T_0 for the binary search tree problem. The corresponding request sequence in $(\mathcal{S}, \mathcal{T})$ will be $(t_{x_1}, \ldots, t_{x_n})$ and the initial state will be any one of the states for which T_0 is the associated tree. Suppose that a $(\mathcal{S}, \mathcal{T})$-algorithm is in state $s = (T, v_r, v_e)$ just before t_{x_i} arrives. Then the algorithm will be forced into a state $s' = (T', v'_r, v'_e)$ for which $v'_r = x_i$ is the request node. The algorithm incurs a cost of $d(s, s')$ for serving t_{x_i}. By definition, this cost is equal to the lowest possible cost for transforming T into T' when the pointer must begin in v_e and must end in v'_e, and when the node containing the key x_i must be visited along the way. This shows that $(\mathcal{S}, \mathcal{T})$ models the dynamic binary search tree problem on the fixed set of N nodes.

We still need to prove that $(\mathcal{S}, \mathcal{T})$ is a LTS. Since $d(s, s') \geq 1$ for all s, s' (due to the initialization cost), we only need to show that d satisfies the triangle inequality. To this end, fix three states $s = (T, v_r, v_e), s' = (T', v'_r, v'_e)$, and $s'' = (T'', v''_r, v''_e)$. Assume that $d(s, s') = t_1$ and $d(s', s'') = t_2$. This means that with a pointer starting at v_e, we can transform T into T', find v'_r and end at v'_e by performing t_1 unit-cost operations. Similarly, with the pointer starting at v'_e, we can transform T' into T'', find v''_r and end at v''_e with t_2 unit-cost operations. We will prove that $d(s, s'') \leq t_1 + t_2$. First, we perform the minimum number of unit-cost operations needed to transform T into T' and end in v'_e. This requires at most t_1 such operations (maybe fewer, since we do not need to find v'_r). After doing so, we perform the minimum number of unit-cost operations needed to transform T' into T'' while also finding v''_r along the way and ending at v''_e. This requires at most t_2 unit-cost operations (in fact, it requires at most $t_2 - 1$ operations since we do not need to perform the initialization operation which is included in $t_2 = d(s', s'')$). Thus, we have transformed T into T'', found v''_r and ended up at v''_e with at most $t_1 + t_2$ unit-cost operations. Hence, $d(s, s'') \leq t_1 + t_2 = d(s, s') + d(s', s'')$ as required.

□

Reordering buffer management. See [86, 1] for a definition of the problem. Note that the ending of a RBM-input is special, since the algorithm has to empty its buffer at the end (even though no new item arrives). Since we need to discuss the repeated version of the problem, we need to specify exactly how this should be done and how to encode into the request sequence that the input ends. The precise way in which this is done is not important, but we need to settle on some definition. We will assume that there is a special item, x_{stop}, which will always be the last item of the input sequence. When a RBM-algorithm tries to load a new item into the buffer and discovers that the only item left is x_{stop}, it will not load this item into the buffer. Instead, the algorithm now knows that no further items will enter the buffer, and so all there is left for the algorithm is to empty its current buffer. Thus, when the item x_{stop} is revealed, an algorithm incurs a cost equal to the minimum number of color switches needed to empty its buffer.

Lemma 16. *The reordering buffer management problem with a buffer of capacity k is Σ-repeatable.*

Proof. We will show that the reordering buffer management problem is Σ-repeatable with parameters $(0, k, 0)$.

Let $\sigma^* = (\sigma_1; \ldots; \sigma_r) \in I^*$ be an instance of the repeated buffer management problem, RBM*. For $1 \leq i \leq r$, the ith-round input sequence has the form $\sigma_i = (x_1, \ldots, x_{n-1}, x_{\text{stop}})$ where x_{stop}

is a special request marking the end of the input stream (note that after x_{stop} arrives, a RBM-algorithm still has to pay some cost in order to empty the buffer). For $i < r$, let $\sigma'_i = \sigma_i - \{x_{\text{stop}}\} = (x_1, \ldots, x_{n-1})$, and let $\sigma'_r = \sigma_r$. The mapping $g : I^* \to I$ is defined by $g(\sigma^*) = \sigma'_1\sigma'_2 \ldots \sigma'_r$. It is clear that g satisfies (25) with $k_1 = 0$ since $|g(\sigma^*)| = |\sigma^*| - (r-1) \leq |\sigma^*|$.

In order to prove (26), fix a deterministic RBM-algorithm ALG (without advice). Let $\sigma^* \in I^*$. We will now define a RBM*-algorithm ALG*: The algorithm ALG* simulates ALG on $g(\sigma^*)$. In each round, the algorithm ALG* makes the exact same color switches (at the same time) as ALG does. When an x_{stop} item is revealed, ALG* pauses its simulation of ALG and empties its current buffer using the minimum number of color switches possible. Recall that x_{stop} is always the last request of a round. When the next round begins, ALG* resumes its simulation of ALG. Note that the set of items in the buffer of ALG* is always a subset of the set of items in the buffer of ALG. Also, ALG* pays a cost of at most k for emptying its buffer at the end of each round. Thus, $\texttt{ALG}^*(\sigma^*) \leq \texttt{ALG}(g(\sigma^*)) + kr$.

We will show that (27) holds. Let $\sigma^* \in I^*$. Consider an optimal sequence of color switches made by $\texttt{OPT}^*_\Sigma$ for serving σ^*. The exact same color switches will also serve all items in $g(\sigma^*)$, since $g(\sigma^*)$ contains the same items in the same order as σ^*. The fact that σ^* contains some additional x_{stop} items which forces $\texttt{OPT}^*_\Sigma$ to empty its buffer at the end of each round is only a disadvantage for $\texttt{OPT}^*_\Sigma$. Thus, $\texttt{OPT}^*_\Sigma(\sigma^*) \geq \texttt{OPT}(g(\sigma^*))$.

□

E Omitted proofs

E.1 Assuming same initial state

We prove Lemma 17 which justifies an assumption regarding initial states made in e.g. Theorem 1. Define $\text{init}(\sigma)$ to be the initial state of the input σ. Let P be an online problem and supose that we are given a probability distribution p such that $\mathbb{E}_p[\texttt{D}(\sigma)] \geq c\,\mathbb{E}_p[\texttt{OPT}(\sigma)] + \alpha$ for every deterministic algorithm D and such that $\text{supp}(p)$ is finite. Let S_p be the set of all initial states appearing in at least one input in $\text{supp}(p)$. For $s \in S_p$, let p_s be the probability distribution obtained from p by conditioning on the event that the initial state of the input is s, that is:

$$p_s(\sigma) = \begin{cases} 0 & \text{if } \text{init}(\sigma) \neq s \\ \frac{p(\sigma)}{\Pr_{\sigma\sim p}[\text{init}(\sigma)=s]} & \text{else}^{12} \end{cases}$$

Lemma 17. *Let p, p_s, and S_p be as defined above. Then there must exist some $s \in S_p$ such that $\mathbb{E}_{p_s}[\texttt{D}(\sigma)] \geq c\,\mathbb{E}_{p_s}[\texttt{OPT}(\sigma)] + \alpha$ for every deterministic* P*-algorithm* D*.*

Proof. Assume by way of contradiction that this is not the case. Then, for every $s \in S_p$, there exists a deterministic algorithm $\texttt{D}_s$ such that $\mathbb{E}_{p_s}[\texttt{D}_s(\sigma)] < c\,\mathbb{E}_{p_s}[\texttt{OPT}(\sigma)] + \alpha$. We now define a P-algorithm $\texttt{D}'$ as follows: Let s be the initial state of the input (which is known to the algorithm before the first request arrives). The algorithm $\texttt{D}'$ uses the algorithm $\texttt{D}_s$ to serve the input. Using the law of total expectation, we get

$$\begin{aligned} \mathbb{E}_{\sigma\sim p}[\texttt{D}'(\sigma) - c\texttt{OPT}(\sigma)] &= \mathbb{E}_{s\in S_p}\left[\mathbb{E}_{\sigma\sim p_s}[\texttt{D}'(\sigma) - c\texttt{OPT}(\sigma)]\right] \\ &= \mathbb{E}_{s\in S_p}\left[\mathbb{E}_{\sigma\sim p_s}[\texttt{D}_s(\sigma) - c\texttt{OPT}(\sigma)]\right] \\ &< \alpha, \end{aligned}$$

[12] Note that since $\text{supp}(p)$ is finite, so is S_p and hence $\Pr_{\sigma\sim p}[\text{init}(\sigma) = s] > 0$ for all $s \in S_p$.

which contradicts the assumption that $\mathbb{E}_{\sigma\sim p}[\mathtt{D}(\sigma) - c\mathtt{OPT}(\sigma)] \geq \alpha$ for every deterministic algorithm $\mathtt{D}$. □

E.2 Proof of Theorem 11

Before giving the proof of Theorem 11, we introduce some terminology and prove a lemma. For $x, y \in [q]^n$, we denote the Hamming distance between x and y by $d(x, y)$.

Definition 15. Let $n, q, r \in \mathbb{N}$ and assume that $q \geq 2$. Let M be a set consisting of strings (codewords) from $[q]^n$. We say that M is a *q-ary anti-covering of radius r and length n* if for every $x \in [q]^n$, there exists a codeword $y \in M$ such that $d(x, y) \geq r$.

Lemma 18. *Let $n, q \in \mathbb{N}$ and assume that $q \geq 2$. Furthermore, let $0 < \alpha < 1/q$. There exists a q-ary anti-covering code of radius $(1-\alpha)n$ and length n of size at most*

$$2^{K_{1/q}(\alpha)n}(n+1)\ln(q)n + 1.$$

Proof. We apply the probabilistic method. Fix $n, q \in \mathbb{N}$ and $0 < \alpha < 1/q$. Let M be a set of m strings chosen uniformly at random and independently from $[q]^n$. Our goal is to show that if m is sufficiently large, then with some non-zero probability the set M is an anti-cover of radius $(1-\alpha)n$.

Fix $x \in [q]^n$. For any $y \in M$, the probability that $d(x, y) \geq (1-\alpha)n$ can be bounded from below using an inverse Chernoff bound ((60)):

$$\Pr[d(x, y) \geq (1-\alpha)n] \geq \frac{2^{-K_{1/q}(\alpha)n}}{n+1} := s.$$

Thus, the probability that there does not exists some $y \in M$ such that $d(x, y) \geq (1-\alpha)n$ is at most

$$\Pr[\forall y \in M : d(x, y) < (1-\alpha)n] \leq (1-s)^m.$$

Applying the union bound over all strings in $[q]^n$ shows that the probability that M is not an anti-covering of radius $(1-\alpha)n$ is at most

$$\begin{aligned}\Pr[M \text{ is not an anti-covering}] &\leq q^n(1-s)^m \\ &\leq q^n e^{-ms} = e^{\ln(q)n - ms}.\end{aligned}$$

Here, we used that $(1-a)^b \leq e^{-ba}$ for $b > 1, a \leq 1$. Set $m = \lfloor s^{-1}\ln(q)n + 1 \rfloor$. Then m is an integer strictly larger than $s^{-1}\ln(q)n$ and hence the probability that M is not an anti-covering of radius r is strictly smaller than 1. Thus, the probability that M is in fact such an anti-covering is strictly positive. This finishes the proof since

$$\begin{aligned}m &= \lfloor s^{-1}\ln(q)n + 1 \rfloor \\ &\leq 2^{K_{1/q}(\alpha)n}(n+1)\ln(q)n + 1.\end{aligned}$$

□

Proof of Theorem 11. Let $x = x_1 \dots x_n$ be the input string. The algorithm computes a q-ary anti-covering code, M^*, of radius $(1-\alpha)n$ and length n with as few codewords as possible (recall that n is revealed to the algorithm before it has to output its first answer). This is simply done by an

exhaustive search. The oracle computes the same anti-covering code M^* as the algorithm. Then, the oracle selects a codeword $y = y_1, \ldots, y_n \in M^*$ such that $d(x, y) \geq (1 - \alpha)n$ and writes the (lexicographical) index of y in M^* onto the advice tape. The algorithm learns the index of y, and thereby also y itself, from the advice. In round i, the algorithm answers y_i. Since $d(x, y) \geq (1-\alpha)n$, the cost incurred by the algorithm will be at most αn.

From Lemma 18, we get that the number of codewords in M^* is at most $2^{K_{1/q}(\alpha)n}(n+1)\ln(q)n+1$. Thus, the number b of advice bits needed to encode the index of a codeword in M^* is at most

$$\begin{aligned} b &\leq \log\left(2^{K_{1/q}(\alpha)n}(n+1)\ln(q)n+1\right) \\ &= K_{1/q}(\alpha)n + \log((n+1)\ln(q)n+1) = K_{1/q}(\alpha)n + O(\log n + \log\log q). \end{aligned}$$

□

F Statement of main results for maximization problems

In what follows, we will state the main results (and definitions) of the paper for the case of Σ-repeatable maximization problems. The proofs are omitted since they can be obtained from the proofs of the corresponding results for minimization with only minor changes. When dealing with a maximization problem, Definition 3 and Definition 4 still applies, except that the word "cost" should be replaced with the word "score". The definition of P^*_Σ (Definition 5) applies as stated to maximization problems. The information-theoretic direct product theorem (Theorem 7) and its application looks as follows for maximization problems.

Theorem 23 (Theorem 7 for maximization problems)**.** *Let* P *be an online maximization problem and let p^r be an r-round input distribution with associated cost-functions* cost_i*. Furthermore, let* `ALG` *be a deterministic algorithm reading at most b bits of advice on every input in the support of p^r.*

Assume that there exists a concave and increasing function $f : [0, \infty] \to \mathbb{R}$ such that for every $w \in \mathcal{W}_i$ where $\Pr[W_i = w] > 0$, the following holds:

$$\mathbb{E}[\mathrm{score}_i(\texttt{ALG})|W_i = w] \leq f\left(D_{KL}(p_{i|w}\|p_i)\right) \tag{64}$$

Then,

$$\mathbb{E}[\mathrm{score}(\texttt{ALG})] \leq rf(b/r). \tag{65}$$

Lemma 19 (Lemma 2 for maximization problems)**.** *Fix $r \geq 1$. Let* P *be a maximization problem. Let $p : I \to [0, 1]$ be an input distribution, where I_s is a finite set of P-inputs with the same initial state s and length at most n'. Assume that for every deterministic* P*-algorithm* `ALG` *without advice, it holds that $\mathbb{E}_{\sigma\sim p}[\texttt{ALG}(\sigma)] \leq t$. Also, let M be the largest score that any* P*-algorithm can achieve on any input from I_s.*

Then, there exists an r-round input distribution, p^r, over P^*_Σ-inputs with at most rn' requests in total, such that any deterministic P^*_Σ-algorithm reading at most b bits of advice (on inputs of length at most rn') has an expected score of at most $r(t + 2M\sqrt{b/r})$. Furthermore, $\mathbb{E}_{\sigma^*\sim p^r}[\texttt{OPT}^*_\Sigma(\sigma^*)] = r\,\mathbb{E}_{\sigma\sim p}[\texttt{OPT}(\sigma)]$.

We need to make some obvious changes to Definition 6 in order to obtain a proper definition of being Σ-repeatable for maximization problems.

Definition 16 (Definition 6 for maximization problems)**.** Let $k_1, k_2, k_3 \geq 0$. An online maximization problem P is Σ-*repeatable with parameters* (k_1, k_2, k_3) if there exists a mapping $g : I^* \to I$ with the following properties:

1. For every $\sigma \in I^*$,
$$|g(\sigma^*)| \leq |\sigma^*| + k_1 r, \tag{66}$$
where r is the number of rounds in σ^*.
2. For any deterministic P-algorithm $\texttt{ALG}$, there is a deterministic P^*_Σ-algorithm $\texttt{ALG}^*$ such that for every $\sigma^* \in I^*$,
$$\texttt{ALG}^*(\sigma^*) \geq \texttt{ALG}(g(\sigma^*)) - k_2 r, \tag{67}$$
where r is the number of rounds in σ^*.
3. For every $\sigma^* \in I^*$,
$$\texttt{OPT}^*_\Sigma(\sigma^*) \leq \texttt{OPT}(g(\sigma^*)) + k_3 r, \tag{68}$$
where r is the number of rounds in σ^*.

Theorem 24 (Theorem 1 for maximization problems)**.** *Let* P *be a* Σ*-repeatable online maximization problem. Suppose that for every* $\varepsilon > 0$ *and every* α*, there exists an input distribution* $p_{\alpha,\varepsilon} : I \to [0,1]$ *with finite support such that* $\mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{OPT}(\sigma)] \geq (c-\varepsilon)\,\mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{D}(\sigma)] + \alpha$ *for every deterministic algorithm* $\texttt{D}$ *without advice. Then, every randomized algorithm reading at most* $o(n)$ *bits of advice on inputs of length* n *has a competitive ratio of at least* c*.*

Definition 17 (Definition 9 for maximization problems)**.** Let P be a maximization problem and let $c > 1$ be such that the expected competitive ratio of every randomized P-algorithm is at least c. We say that P is *compact* if for every $\varepsilon > 0$ and every $\alpha \geq 0$, there exists an input distribution $p_{\alpha,\varepsilon}$ such that

1. If $\texttt{ALG}$ is a deterministic online algorithm (without advice) then
$$\mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{OPT}(\sigma)] \geq (c-\varepsilon) \cdot \mathbb{E}_{p_{\alpha,\varepsilon}}[\texttt{ALG}(\sigma)] + \alpha. \tag{69}$$
2. The support $\mathrm{supp}(p_{\alpha,\varepsilon})$ of $p_{\alpha,\varepsilon}$ is finite.

The maximization version of Theorem 2 follows easily from Definition 17 and Theorem 24 (and Corollary 2).

Theorem 25 (Theorem 2 for maximization problems)**.** *Let* P *be compact and* Σ*-repeatable maximization problem with at most* $2^{n^{O(1)}}$ *inputs of length* n*, and let* c *be a constant not depending on* n*. The following are equivalent:*

1. *For every* $\varepsilon > 0$*, there exists a randomized* $(c+\varepsilon)$*-competitive* P*-algorithm without advice.*
2. *For every* $\varepsilon > 0$*, there exists a deterministic* $(c+\varepsilon)$*-competitive* P*-algorithm with advice complexity* $o(n)$*.*